



**HAL**  
open science

# An efficient joint model for high dimensional longitudinal and survival data via generic association features

Van Tuan Nguyen, Adeline Fermanian, Antoine Barbieri, Sarah Zohar, Anne-Sophie Jannot, Simon Bussy, Agathe Guilloux

► **To cite this version:**

Van Tuan Nguyen, Adeline Fermanian, Antoine Barbieri, Sarah Zohar, Anne-Sophie Jannot, et al.. An efficient joint model for high dimensional longitudinal and survival data via generic association features. *Biometrics*, 2024, 80 (4), 10.1093/biomtc/ujae149 . hal-04875977

**HAL Id: hal-04875977**

**<https://hal.science/hal-04875977v1>**

Submitted on 12 Feb 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

## An efficient joint model for high dimensional longitudinal and survival data via generic association features

Van Tuan Nguyen<sup>1,2,\*</sup>, Adeline Fermanian<sup>2</sup>, Antoine Barbieri<sup>3</sup>,

Sarah Zohar<sup>4,5</sup>, Anne-Sophie Jannot<sup>4,5,6</sup>, Simon Bussy<sup>1</sup>, and Agathe Guilloux<sup>4,5</sup>

<sup>1</sup>LOPF, Califrais’ Machine Learning Lab, Paris, France, <sup>2</sup>Laboratoire de Probabilités, Statistique et Modélisation, LPSM, Univ. Paris Cité, F-75005, Paris, France, <sup>3</sup>Univ. Bordeaux, INSERM, U1219, France,

<sup>4</sup>INSERM, Centre de recherche des Cordeliers, Univ. Sorbonne, Univ. Paris Cité, F-75006, Paris, France,

<sup>5</sup>HeKA, Inria Paris, F-75015, Paris, France, <sup>6</sup>French National Rare Disease Registry (BNDMR),

Greater Paris University Hospitals (AP-HP), Paris, France.

\**email*: vantuan.nguyen@califrais.fr

**SUMMARY:** This paper introduces a prognostic method called FLASH that addresses the problem of joint modelling of longitudinal data and censored durations when a large number of both longitudinal and time-independent features are available. In the literature, standard joint models are either of the shared random effect or joint latent class type. Combining ideas from both worlds and using appropriate regularisation techniques, we define a new model with the ability to automatically identify significant prognostic longitudinal features in a high-dimensional context, which is of increasing importance in many areas such as personalised medicine or churn prediction. We develop an estimation methodology based on the EM algorithm and provide an efficient implementation. The statistical performance of the method is demonstrated both in extensive Monte Carlo simulation studies and on publicly available medical datasets. Our method significantly outperforms the state-of-the-art joint models in terms of C-index in a so-called “real-time” prediction setting, with a computational speed that is orders of magnitude faster than competing methods. In addition, our model automatically identifies significant features that are relevant from a practical point of view, making it interpretable, which is of the greatest importance for a prognostic algorithm in healthcare.

**KEY WORDS:** Joint models; Longitudinal data; High-dimensional statistics; Survival analysis


### 1. Introduction

In healthcare, it is increasingly common to record the values of longitudinal features (e.g. biomarkers such as heart rate or haemoglobin level) up to the occurrence of an event of interest for a subject, such as rehospitalisation, relapse, or disease progression. Moreover, in large observational databases such as claims databases, with electronic health records, the amount of recorded data per patient is often very large and growing over time. However, there is currently no tool that can simultaneously process high-dimensional longitudinal signals and perform real-time predictions, i.e. give predictions at any time after only one estimation/training step. While landmark approaches (see, e.g., Devaux et al., 2022) can handle high-dimensional longitudinal features, they require a separate training for each prediction time. An alternative is to use “joint modeling” to handle longitudinal and survival outcomes together.

The latter has received considerable attention in the last two decades (Tsiatis and Davidian, 2004; Rizopoulos and Ghosh, 2011; Hickey et al., 2016). More specifically, it consists of defining (*i*) a time-to-event model, (*ii*) a longitudinal marker model, and (*iii*) linking both models via a common latent structure. Numerical studies suggest that these approaches are among the most satisfactory for incorporating all longitudinal information into a survival model (Yu et al., 2004), and are better than landmark approaches, which use only information from individuals at risk at the landmark time (see, e.g., Devaux et al., 2022). They have the additional advantage of making more efficient use of the data as information on survival is also used to model the longitudinal markers. More importantly, they require only one training, regardless of the number of prediction times.

There are two main approaches to linking longitudinal and survival models. On the one hand, in shared parameter joint models (SREMs), characteristics of the longitudinal marker, typically some random effects learned in a linear mixed model, are included as covariates in the survival model (Wulfsohn and Tsiatis, 1997; Andrinopoulou and Rizopoulos, 2016). On the other hand, joint latent class models (JLCMs), inspired by mixture-of-experts modelling (Masoudnia and Ebrahimpour, 2014), assume that the dependence between the time-to-event and the longitudinal marker is fully captured by a latent class structure (Lin et al., 2002; Proust-Lima et al., 2014), which amounts to assuming that the population is heterogeneous and that there are

homogeneous latent classes that share the same marker trajectories and prognosis. JLCMs offer a computationally attractive alternative to SREMs, especially in a high-dimensional context. These two models are illustrated in Figure 1.




**Figure 1:** Graphical representation of SREMs (left) and JLCMs (right). The variable  $X$  represents time-independent features,  $Y$  the longitudinal markers,  $T$  the time-to-event, and  $G$  the latent class membership.

Unfortunately, joint models have predominantly focused on univariate longitudinal markers (Andrinopoulou et al., 2020). To adapt such models to a multivariate setting, the common approach is to fit multiple univariate joint models separately to each longitudinal marker (Wang et al., 2012), which does not account for interactions between longitudinal markers (Jaffa et al., 2014; Kang and Song, 2022; Lin et al., 2002). Furthermore, issues arising from the high-dimensional context—e.g. computational power, limits of numerical estimation—have, to our knowledge, never been considered in the analyses, and the number of longitudinal markers considered in numerical studies remains very low, typically up to 5 (Hickey et al., 2016; Murray and Philipson, 2022; Rustand et al., 2024).

The aim of this article is to propose a new joint model called FLASH (Fast joint model for Longitudinal And Survival data in High dimension), together with an efficient inference methodology. The model is inspired by both JLCMs and SREMs, but is designed to scale to high-dimensional longitudinal markers. The general idea is to use generic features extracted from the longitudinal markers directly in the survival model and to use regularization in an Expectation-Maximization (EM) algorithm. The main difference with SREMs is that these features, called association features, are assumed to be independent of the modeling assumptions in the longitudinal submodel. As a result, the model is very efficient to train—the likelihood is closed-form in a Gaussian setting and does not require Monte Carlo approximations, as is often the case with SREMs, making it suitable for high-dimensional longitudinal markers. Moreover, it allows the use of high-dimensional and generic feature extraction functions that characterize longitudinal markers, rather than just the random effects, resulting in a model that is generic enough to be adapted to different types of longitudinal markers and prior information on the problem. The use of elastic net and sparse group lasso regularization enables automatic selection of relevant association features, resulting in an interpretable model that retains only the significant longitudinal markers.

Finally, our model allows us to define a “real-time” prediction methodology where, once the parameters of the model have been learned, we can compute a predictive marker that, given only the time-independent and longitudinal features up to a given point in time, outputs a risk for each subject at that point in time. It differs from traditional approaches (Proust-Lima et al., 2014) that require knowledge of the survival labels, which are unknown in the “real-time” prediction setting.



**Figure 2:** Graphical representation of the FLASH model. The variable  $X$  represents time-independent features,  $Y$  the longitudinal markers,  $T$  the time-to-event and  $G$  the latent class membership of an individual.

In summary, we introduce a new method for predicting survival risk with high-dimensional longitudinal features that is both interpretable and computationally efficient, thus providing a powerful tool for real-time clinical decision making, for example in patient monitoring.

A precise description of the model is given in Section 2. In Section 3, we present our assumptions and inference methodology based on maximizing the likelihood with a regularized variant of the EM algorithm. Section 4 introduces our evaluation methodology and the competing methods. In Section 5, we apply our method to datasets from two simulation studies and three publicly available medical datasets (PBCseq, AIDS, Sepsis). We show that FLASH outperforms the competitors and

selects relevant features from a medical perspective. Finally, we discuss the obtained results in Section 6. The code to implement the model and reproduce the experiments is publicly available at <https://github.com/Califrais/flash.git>.

## 2. Model

In this section we describe the FLASH model, which consists of three sub-models: a multinomial logistic regression defining the probability of belonging to a latent class, a generalized linear mixed model for each latent class describing the evolution of the longitudinal markers, and finally a Cox class-specific survival model. In all the following, we consider a set of  $n$  independent and identically distributed (i.i.d.) subjects. For each subject  $i \in \{1, \dots, n\}$  we are given some longitudinal markers  $Y_i$ , time-independent features  $X_i \in \mathbb{R}^p$ , a right-censored time-to-event  $T_i \in \mathbb{R}^+$ , and a censoring variable  $\Delta_i \in \{0, 1\}$ .

### 2.1 Latent class membership

We assume that the population consists of  $K \in \mathbb{N}^*$  latent groups. To each subject  $i \in \{1, \dots, n\}$ , we associate a categorical latent variable  $g_i \in \{1, \dots, K\}$ , which encodes its latent class membership. Then, denoting by  $X_i \in \mathbb{R}^p$  the  $p$ -dimensional vector of time-independent features, the latent class membership probability is assumed to take the multinomial logistic regression form (Böhning, 1992), for any  $k \in \{1, \dots, K\}$ ,

$$\mathbb{P}(g_i = k) = \frac{e^{X_i^\top \xi_k}}{\sum_{j=1}^K e^{X_i^\top \xi_j}}, \quad (1)$$

where  $\xi_k \in \mathbb{R}^p$  denotes a vector of coefficients for class  $k$ . This submodel is similar to JLCMs or the C-mix model (Bussy et al., 2019) and assumes that latent-class membership depends only on time-independent features, with the vector  $\xi_k$  quantifying the effect of each time-independent feature in  $X_i$  on the probability that subject  $i$  belongs to the  $k$ -th latent class. The optimal number of latent classes  $K$  can be selected with the Bayesian information criterion (BIC) (Hastie et al., 2009), see Section C.8 of the Supplementary Materials for more details. Note that modeling the probability of latent class membership using multinomial logistic regression is a common assumption in joint models (Lin et al., 2002; Proust-Lima et al., 2014). However, other models such as hidden Markov models can also be used to represent this probability (Bartolucci and Farcomeni, 2015, 2019). We assess the sensitivity to this choice by some experiments in Section C.9 of the Supplementary Materials.

### 2.2 Class-specific longitudinal model

For each subject  $i \in \{1, \dots, n\}$ , we are given  $L \in \mathbb{N}^*$  longitudinal markers. We let, for any  $\ell \in \{1, \dots, L\}$ ,  $Y_i^\ell = (y_i^\ell(t_{i1}^\ell), \dots, y_i^\ell(t_{in_i^\ell}^\ell))^\top \in \mathbb{R}^{n_i^\ell}$  be the vector of repeated measures of a theoretical longitudinal marker  $y_i^\ell$  at observation times (or follow-up visits)  $0 \leq t_{i1}^\ell < \dots < t_{in_i^\ell}^\ell$ . Note that the observation times  $t_{ij}^\ell$ ,  $j = 1, \dots, n_i^\ell$ , can differ between subjects as well as between longitudinal markers, which makes the assumptions on the sampling mechanism very weak. In particular, this setting encapsulates many scenarios considered as missing data, where one individual is not measured while another one is, or where the longitudinal marker of one individual is missing, which here both simply correspond to removing some timestamps from the corresponding grids.

We assume a class-specific generalized linear mixed model (GLMM) for each longitudinal marker, which is a classical model for longitudinal data (Fitzmaurice et al., 2012; Hickey et al., 2016). The GLMM is chosen according to the nature of the markers: Gaussian linear model for continuous markers, logistic regression for a binary marker, Poisson regression for counts. For the continuous markers, given a latent class  $g_i = k$ , for the  $\ell$ -th marker at time  $t \in \{t_{i1}^\ell, \dots, t_{in_i^\ell}^\ell\}$ , we then have

$$y_i^\ell(t_{ij}^\ell) | b_i^\ell, g_i = k \sim \mathcal{N}(m_{ik}^\ell(t_{ij}^\ell), \phi_\ell), \quad (2)$$

where the variance  $\phi_\ell \in \mathbb{R}^+$  is an estimated parameter and the mean  $m_{ik}^\ell$  is defined by

$$m_{ik}^\ell(t) = u^\ell(t)^\top \beta_k^\ell + v^\ell(t)^\top b_i^\ell,$$

where  $u^\ell(t) \in \mathbb{R}^{q_\ell}$  is a row vector of time-varying features with corresponding unknown fixed effect parameters  $\beta_k^\ell \in \mathbb{R}^{q_\ell}$ , and  $v^\ell(t) \in \mathbb{R}^{r_\ell}$  is a row vector of time-varying features with corresponding random effect  $b_i^\ell \in \mathbb{R}^{r_\ell}$ . Flexible representations for  $u^\ell(t)$  can be considered using a vector of time monomials  $u^\ell(t) = (1, t, t^2, \dots, t^\alpha)^\top$ , with  $\alpha \in \mathbb{N}^+$ . We use  $\alpha = 1$  in all our experiments but higher orders could be used. We also let  $v^\ell(t) = (1, t)^\top$ .

Classically, the random effects component is assumed to follow a zero-mean multivariate normal distribution, that is,  $b_i^\ell \sim \mathcal{N}(0, D^{\ell\ell})$  with  $D^{\ell\ell} \in \mathbb{R}^{r_\ell \times r_\ell}$  the variance-covariance matrix. To account for the dependence between the different longitudinal markers, we let  $\text{Cov}[b_i^\ell, b_i^{\ell'}] = D^{\ell\ell'}$  for  $\ell \neq \ell'$ , where  $\text{Cov}[\cdot, \cdot]$  denotes the covariance matrix of two random vectors, and we denote by  $D = (D^{\ell\ell'})_{1 \leq \ell, \ell' \leq L}$  the global variance-covariance matrix which is common to all latent classes. Note that this variance-covariance matrix  $D$  can be easily extended to be class-specific. We assume that all dependencies between longitudinal markers are encapsulated in this matrix  $D$ , which is summarized in the following assumption.

ASSUMPTION 1: For any  $\ell \in \{1, \dots, L\}$  and any  $i \in \{1, \dots, n\}$ , the longitudinal markers  $Y_i^\ell$  are pairwise independent conditionally on  $b_i^\ell$  and  $g_i$ .

This is a standard modelling assumption in joint models (see, e.g., Tsiatis and Davidian, 2004). Then, if we concatenate all longitudinal measurements and random effects of subject  $i$  in, respectively,  $Y_i = (Y_i^{1\top} \dots Y_i^{L\top})^\top \in \mathbb{R}^{n_i}$  and  $b_i = (b_i^{1\top} \dots b_i^{L\top})^\top \in \mathbb{R}^r$  with  $n_i = \sum_{\ell=1}^L n_i^\ell$  and  $r = \sum_{\ell=1}^L r_\ell$ , a consequence of Assumption 1 and Equation (2) is that

$$Y_i | b_i, g_i = k \sim \mathcal{N}(M_{ik}, \Sigma_i), \quad (3)$$

where  $M_{ik} = (m_{ik}^1(t_{i1}^1), \dots, m_{ik}^1(t_{in_i^1}^1), \dots, m_{ik}^L(t_{i1}^L), \dots, m_{ik}^L(t_{in_i^L}^L))^\top \in \mathbb{R}^{n_i}$  and  $\Sigma_i$  is the diagonal matrix whose diagonal is  $(\phi_1 \mathbf{1}_{n_i^1}^\top, \dots, \phi_L \mathbf{1}_{n_i^L}^\top)^\top \in \mathbb{R}^{n_i}$  where  $\mathbf{1}_m$  denotes the vector of  $\mathbb{R}^m$  having all coordinates equal to one.

To extend the GLMM to other types of longitudinal markers (for example, binary or count data), the key point is that the distributions of  $b_i$  and of  $Y_i | b_i, g_i = k$  should be conjugate, so that the likelihood is tractable. It is possible to extend the model to non-conjugate distributions via numerical integration methods but this is not immediate. A detailed discussion of this is given in Section A.1 of the Supplementary Materials. For simplicity, we restrict ourselves to the Gaussian case in this article.

### 2.3 Class-specific Cox survival model

We place ourselves in a classical survival analysis framework. Let the non-negative random variables  $T_i^*$  and  $C_i$  be the time to the event of interest and the censoring time, respectively. We then define the observed time  $T_i = T_i^* \wedge C_i$  and censoring indicator  $\Delta_i = \mathbb{1}_{\{T_i^* \leq C_i\}}$ , where  $a \wedge b$  denotes the minimum between two real numbers  $a$  and  $b$ , and  $\mathbb{1}_{\{\cdot\}}$  is the indicator function which takes the value 1 if the condition in  $\{\cdot\}$  is satisfied, and 0 otherwise. We denote by  $\mathcal{Y}_i^\ell(t^-) = (y_i^\ell(t_{i1}^\ell), \dots, y_i^\ell(t_{iu}^\ell))_{0 \leq t_{iu}^\ell < t}$  the subset of  $Y_i^\ell$  formed from observations up to time  $t$  and by  $\mathcal{Y}_i(t^-)$  the concatenation of the history of all observed longitudinal markers up to  $t$ . Then we consider  $M \in \mathbb{N}^+$  user-defined feature extraction functions  $\Psi_m : \mathcal{Y}_i^\ell(t^-) \rightarrow \Psi_m(\mathcal{Y}_i^\ell(t^-)) \in \mathbb{R}$ ,  $m \in \{1, \dots, M\}$ , which characterise the longitudinal markers. The set of features  $(\Psi_m(\mathcal{Y}_i^\ell(t^-)))_{1 \leq m \leq M}$  should be rich enough to capture all dependencies between longitudinal markers and time-to-event, and is discussed in more detail below. To quantify the effect of the longitudinal markers on time-to-event, we then use an extension of the Cox relative risk model (Cox, 1972), which allows time-varying covariates and was firstly introduced in Andersen and Gill (1982). Note that this model does not fulfill the classical Cox model's assumption of a constant proportional hazard over time. The hazard function in this model takes the form

$$\lambda(t | \mathcal{Y}_i(t^-), g_i = k) = \lambda_0(t) \exp \left( \sum_{\ell=1}^L \sum_{m=1}^M \Psi_m(\mathcal{Y}_i^\ell(t^-)) \gamma_{k,m}^\ell \right) = \lambda_0(t) \exp(\psi_i(t)^\top \gamma_k), \quad (4)$$

where  $\lambda_0$  is an unspecified baseline hazard function that does not depend on  $k$ ,  $\gamma_{k,m}^\ell \in \mathbb{R}$  the joint association parameters, which are the only class-specific objects in this model. We concatenate them in  $\gamma_k = (\gamma_{k,1}^1, \dots, \gamma_{k,M}^1, \dots, \gamma_{k,1}^L, \dots, \gamma_{k,M}^L)^\top \in \mathbb{R}^{LM}$  and define  $\psi_i(t) = (\Psi_1(\mathcal{Y}_i^1(t^-)), \dots, \Psi_M(\mathcal{Y}_i^1(t^-)), \dots, \Psi_1(\mathcal{Y}_i^L(t^-)), \dots, \Psi_M(\mathcal{Y}_i^L(t^-)))^\top \in \mathbb{R}^{LM}$ .

This model can be viewed as a generalization of SREMs (Lin et al., 2002; Rizopoulos and Ghosh, 2011), which have hazard functions of the form  $\lambda_0(t) \exp(\sum_{\ell=1}^L \phi(b_i^\ell, t)^\top \gamma^\ell)$ , where the association between the longitudinal and survival models is captured by the random effects  $b_i^\ell$ . The key idea of our model lies in the following assumption.


ASSUMPTION 2: For any time  $t \geq 0$ , the hazard rate at time  $t$  conditionally on the history of the longitudinal markers up to  $t^-$  and  $g_i$  is independent of  $b_i$ .

Any feature extraction function  $\Psi_m$  of the longitudinal markers can be considered in the hazard function. Simple examples of such functions are the maximum or the sum (or the mean) of the longitudinal features, respectively defined by

$$\Psi_m(\mathcal{Y}_i^\ell(t^-)) = \max_{j:t_{ij}^\ell < t} \{y_i^\ell(t_{i1}^\ell), \dots, y_i^\ell(t_{ij}^\ell)\}, \quad \Psi_m(\mathcal{Y}_i^\ell(t^-)) = \sum_{j:t_{ij}^\ell < t} y_i^\ell(t_{ij}^\ell).$$

For illustration, Figure 3 shows how these association features change with time. In practice, our rationale is to use a variety of feature extraction functions  $\Psi_m$ , such as absolute energy over time, statistics on autocorrelation, or Fourier and wavelet basis projections, and then perform feature selection via regularisation to learn which ones are predictive for the underlying task. This will be described in more detail in Section 3.2. Note that a crucial aspect of this model is that the extracted vector, also called extracted features or association features,  $\psi_i(t)$ , does not depend on the modelling assumptions in the longitudinal submodel of Subsection 2.2 — that is, does not depend on  $b_i$  other than through the history  $\mathcal{Y}_i(t^-)$ .

The FLASH model is summarised in Figure 2 which shows that our model is a combination of SREMs and JLCMs where both random effects and latent classes account for the dependence between longitudinal markers and time-to-event.



**Figure 3:** On the **top**, observed longitudinal markers of 2 individuals from the FLASH\_simu dataset. On the **bottom**, association features applied on the observed longitudinal markers: the maximum (**left**), the mean (**center**), and the sum (**right**).

### 3. Inference

Now that we have introduced all the components of our model, in this section we derive the form of its likelihood, present the regularisation strategy that deals with the high dimensionality of the data, and finally present our variant of the EM algorithm used to minimise the penalised negative log-likelihood.

#### 3.1 Likelihood

Consider a training cohort of  $n$  i.i.d. subjects  $\mathcal{D}_n = ((X_1, Y_1, T_1, \Delta_1), \dots, (X_n, Y_n, T_n, \Delta_n))$ . For simplicity, we slightly abuse notation and use the same notation  $f^*$  for the true (joint) density or probability mass function of the various random variables in our model. Similarly, we denote by  $f_\theta$  the candidates for estimating the densities  $f^*$  that satisfy the model assumptions of Section 2, where we have concatenated in  $\theta$  all  $P \in \mathbb{N}^+$  unknown parameters:

$$\theta = (\xi_1^\top, \dots, \xi_K^\top, \beta_1^\top, \dots, \beta_K^\top, \phi^\top, D, \lambda_0(\tau_1), \dots, \lambda_0(\tau_J), \gamma_1^\top, \dots, \gamma_K^\top)^\top \in \mathbb{R}^P,$$

where  $\beta_k = (\beta_k^1 \dots \beta_k^\ell)^\top \in \mathbb{R}^q$  with  $q = \sum_{\ell=1}^L q_\ell$  for any  $k \in \{1, \dots, K\}$ ,  $\phi = (\phi_1, \dots, \phi_L)^\top$  and where we use the vectorization of the matrix  $D$  although this is not written explicitly. Note that we classically (see, e.g., Klein, 1992) estimate  $\lambda_0$  by a function taking mass at each failure time  $\tau_j \in (\tau_1, \dots, \tau_J)$ , where  $(\tau_1, \dots, \tau_J)$  denote the  $J \in \mathbb{N}^+$  unique failure times (obtained from  $(T_1, \dots, T_n)$  removing the duplicates and keeping only the uncensored times  $T_i$  for which  $\Delta_i = 1$ ). In this way, the estimation of the function  $\lambda_0$  amounts to the estimation of the vector  $(\lambda_0(\tau_1), \dots, \lambda_0(\tau_J))$ .

First, conditioning on the latent classes, we have

$$f^*(T_i, \Delta_i, Y_i) = \sum_{k=1}^K f^*(g_i = k) f^*(T_i, \Delta_i | Y_i, g_i = k) f^*(Y_i | g_i = k).$$

This yields the negative log-likelihood

$$\mathcal{L}_n(\theta) = -n^{-1} \sum_{i=1}^n \log \sum_{k=1}^K f_\theta(g_i = k) f_\theta(T_i, \Delta_i | Y_i, g_i = k) f_\theta(Y_i | g_i = k). \quad (5)$$

Assuming that both the censoring mechanism and the stochastic mechanism generating the observation times of the longitudinal markers are non-informative (Rizopoulos and Ghosh, 2011), the joint density of  $(T_i, \Delta_i)$  can be factorized into a part depending on the distribution of  $T_i^*$  and a part depending on that of  $C_i$ , so that

$$\begin{aligned} f^*(T_i, \Delta_i | Y_i, g_i = k) &\propto f^*(T_i | Y_i, g_i = k)^{\Delta_i} S^*(T_i | Y_i, g_i = k)^{1-\Delta_i} \\ &= \lambda^*(T_i | Y_i, g_i = k)^{\Delta_i} S^*(T_i | Y_i, g_i = k), \end{aligned} \quad (6)$$

where  $S^*$  and  $\lambda^*$  are the survival and hazard function associated with the density  $f^*$  of  $T_i^*$ .

Under the assumptions given in the previous subsections, all terms in (5) can be computed in closed form. Indeed,  $f_\theta(g_i = k)$  is given by (1) and the density function  $f_\theta(Y_i | g_i = k)$  can be derived from the distribution of  $b_i$  and (3) (detailed calculations are given in Section A of the Supplementary Materials). Furthermore, following Equation (6), we have  $f_\theta(T_i, \Delta_i | Y_i, g_i = k) \propto \lambda(T_i | \mathcal{Y}_i(T_i^-, g_i = k))^{\Delta_i} S_k(T_i)$ , where  $S_k(t) = \exp\left(-\int_0^t \lambda(s | \mathcal{Y}_i(s^-, g_i = k)) ds\right)$  is the survival function of subject  $i$  given

that it belongs to latent class  $k$ . Since the baseline hazard function  $\lambda_0$  takes mass only at each failure time  $\tau_j \in (\tau_1, \dots, \tau_J)$  then the integration over the survival process  $S_k(t)$  is simply a finite sum over the process evaluated at the  $J$  failure times. Then, we rewrite the function  $S_k$ , for any  $t \geq 0$ , as

$$S_k(t) = \exp\left(-\sum_{j=1}^J \lambda(\tau_j) \mathbb{1}_{\{\tau_j \leq t\}}\right) = \exp\left(-\sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k) \mathbb{1}_{\{\tau_j \leq t\}}\right).$$

The fact that  $f_\theta(T_i, \Delta_i | Y_i, g_i = k)$  is closed-form is one of the major advantages of our model over standard SREMs. Indeed, computing this density in SREMs usually requires integrating it with respect to the distribution of the random effects  $b_i$ , leading to intractable integrals in the log-likelihood function. These integrals are typically estimated using Monte Carlo techniques (Hickey et al., 2018), which are computationally intensive and require additional assumptions on the allowed association functions  $\psi_i$ . These approaches usually do not scale in a high-dimensional context.

To minimize (5) with respect to  $\theta$ , we use the EM algorithm, which is the common choice in the literature (Wulfsohn and Tsiatis, 1997; Lin et al., 2002). This requires deriving what we call the negative “complete” log-likelihood, that is, an estimation of the joint density  $f^*(T_i, \Delta_i, Y_i, b_i, g_i)$ , where the random effect  $b_i$  and the latent class  $g_i$  are not observed. To this end, we need the following independence assumption.

**ASSUMPTION 3:** For any  $i \in \{1, \dots, n\}$  and  $\ell \in \{1, \dots, L\}$ , the random effects  $b_i^\ell$  are independent of the latent class membership  $g_i$ , and remain independent of it conditionally on  $T_i$ ,  $\Delta_i$ , and  $Y_i$ .

This assumption states that subject-and-longitudinal marker specific random effects  $b_i^\ell$  do not depend on the latent class membership. Then, we have

$$\begin{aligned} f^*(T_i, \Delta_i, Y_i, b_i, g_i) &= f^*(b_i, g_i) f^*(Y_i | b_i, g_i) f^*(T_i, \Delta_i | Y_i, b_i, g_i) \\ &= f^*(b_i, g_i) f^*(Y_i | b_i, g_i) f^*(T_i, \Delta_i | Y_i, g_i) && \text{(by Assumption 2)} \\ &= f^*(b_i) f^*(g_i) f^*(Y_i | b_i, g_i) f^*(T_i, \Delta_i | Y_i, g_i). && \text{(by Assumption 3)} \end{aligned}$$

The negative complete log-likelihood is then given by

$$\begin{aligned} \mathcal{L}_n^{\text{comp}}(\theta) &= -n^{-1} \sum_{i=1}^n \left( \log f_\theta(b_i) + \sum_{k=1}^K \mathbb{1}_{\{g_i=k\}} (\log \mathbb{P}_\theta(g_i = k) + \log f_\theta(Y_i | b_i, g_i = k)) \right. \\ &\quad \left. + \log f_\theta(T_i, \Delta_i | Y_i, g_i = k) \right), \end{aligned} \quad (7)$$

where  $f_\theta(b_i)$  is the density of a multivariate gaussian  $\mathcal{N}(0, D)$  distribution and  $f_\theta(Y_i | b_i, g_i = k)$  is typically the density of a  $\mathcal{N}(M_{ik}, \Sigma_i)$  distribution.

### 3.2 Penalized objective

To avoid overfitting and provide interpretation on which longitudinal markers are relevant for predicting time-to-event, we propose to minimize the penalized negative log-likelihood

$$\mathcal{L}_n^{\text{pen}}(\theta) = \mathcal{L}_n(\theta) + \Omega(\theta) = \mathcal{L}_n(\theta) + \sum_{k=1}^K \zeta_{1,k} \Omega_1(\xi_k) + \sum_{k=1}^K \zeta_{2,k} \Omega_2(\gamma_k), \quad (8)$$

where  $\Omega_1$  is an elastic net regularization (Zou and Hastie, 2005),  $\Omega_2$  is a sparse group lasso regularization (Simon et al., 2013), and  $(\zeta_{1,k}, \zeta_{2,k})^\top \in (\mathbb{R}^+)^2$  regularization hyperparameters that need to be tuned. An advantage of this regularisation strategy is its ability to perform feature selection and to identify the most important features (longitudinal markers and time-independent) relative to the prediction objective. On the one hand, the support of  $\xi_k$ , controlled by the  $\ell_1$  term in  $\Omega_1$ , provides information about the time-independent features involved in the  $k$ -th latent class membership while the  $\ell_2$  regularization allows to handle correlations between time-independent features. On the other hand, for the sparse group lasso penalty, a group  $\ell$  corresponds to a trajectory, i.e. a longitudinal marker. Thus, if  $\gamma_k^\ell$  is completely zero (thanks to the group lasso part), it means that the  $\ell$ -th longitudinal process is discarded by the model in terms of risk effect for the  $k$ -th latent class. Then, the sparse part of the penalty allows a selection of association features for each trajectory: for  $\gamma_k^\ell$  that are not completely zeroed, their support informs about the association features involved in the risk of the  $k$ -th latent class event for the  $\ell$ -th longitudinal marker.

### 3.3 Optimization

Given our regularization strategy, we employ an extended version of the EM algorithm (McLachlan and Krishnan, 2007) which we now briefly outline. Extensive details on the algorithm are given in Section A of the Supplementary Materials.

Our final optimization problem writes

$$\hat{\theta} \in \underset{\theta \in \mathbb{R}^P}{\text{argmin}} \mathcal{L}_n^{\text{pen}}(\theta). \quad (9)$$

Assume that we are at step  $w + 1$  of the EM algorithm, with current iterate denoted by  $\theta^{(w)}$ , then the algorithm consists in the following two steps:

- E-step: compute the expected negative complete log-likelihood conditional on the current estimate of the parameters  $\theta^{(w)}$ , that is,  $\mathcal{Q}_n(\theta, \theta^{(w)}) = \mathbb{E}_{\theta^{(w)}}[\mathcal{L}_n^{\text{comp}}(\theta) | \mathcal{D}_n]$ .
- M-step: find  $\theta^{(w+1)} \in \underset{\theta \in \mathbb{R}^P}{\text{argmin}} \mathcal{Q}_n^{\text{pen}}(\theta, \theta^{(w)})$ , where  $\mathcal{Q}_n^{\text{pen}}(\theta, \theta^{(w)}) = \mathcal{Q}_n(\theta, \theta^{(w)}) + \Omega(\theta)$  and  $\Omega(\theta)$  is the penalization defined in (8).

Under our assumptions, we can show that computing  $\mathcal{Q}_n(\theta, \theta^{(w)})$  reduces to computing the expectations  $\mathbb{E}_{\theta^{(w)}}[b_i | T_i, \Delta_i, Y_i]$  and  $\mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i]$ , and the probabilities  $\mathbb{P}_{\theta^{(w)}}[g_i = k | T_i, \Delta_i, Y_i]$ ,  $k \in \{1, \dots, K\}$ , see Section A.1 of the Supplementary Materials for their exact expressions.

Concerning the M-step, we divide the problem into several updates for which we minimize  $\mathcal{Q}_n^{\text{pen}}(\theta, \theta^{(w)})$  with respect to blocks of coordinates of  $\theta$  separately. The order of these updates matters. The updates for  $D$ ,  $(\beta_k)_{k \in \{1, \dots, K\}}$ ,  $\lambda_0$ , and  $\phi$  are easily obtained in closed-form. The update for  $\xi_k^{(w)}$  reduces to the non-smooth convex minimization problem

$$\xi_k^{(w+1)} \in \underset{\xi \in \mathbb{R}^P}{\text{argmin}} \mathcal{F}_{1,k}(\xi) + \zeta_{1,k} \Omega_1(\xi), \quad (10)$$

where  $\mathcal{F}_{1,k}$  is a convex function with respect to  $\xi$ . Problem (10) is then solved using a quasi-Newton method, the L-BFGS-B algorithm (Zhu et al., 1997). The update for  $\gamma_k^{(w)}$  has a similar expression and is solved using proximal gradient descent (Boyd and Vandenberghe, 2004). We refer to Section A of the Supplementary Materials for all details and proofs. The final algorithm is also given with a discussion of its convergence properties.

#### 4. Evaluation methodology

In this section, we present our evaluation strategy to assess real-time prediction performance of our model and briefly introduce the models used for comparison.

##### 4.1 Real-time prediction and evaluation strategy

Developments in joint models have focused primarily on modeling and estimation, and most studies do not consider goodness-of-fit or predictive performance of latent class membership or time-to-event (Hickey et al., 2016). However, for real-time or daily predictions, practitioners need predictive prognostic tools to evaluate and compare survival models. Therefore, we place ourselves in a so-called “real-time” prediction setting. Once the learning phase for the model has been completed on a training set, so that one obtains  $\hat{\theta}$  from (9) using the approach described in Section 3.3, we want to make real-time predictions. More precisely, for each subject  $i$ , we seek to provide a predictive marker, typically the probability of belonging to a latent class at any time  $t$ , using all the data available up to that time, but without using the supervision labels  $(T_i, \Delta_i)$ , which are a priori not available at any time  $t$ .

**4.1.1 Predictive marker.** In our setting, since each latent class represents the different risk levels of a subject, we choose the probability of latent class membership as the predictive marker. This is similar to what is classically done in JLCMs, where  $\tilde{\pi}_{ik}^{\hat{\theta}} = \mathbb{P}_{\hat{\theta}}[g_i = k | T_i, \Delta_i, Y_i]$  is typically used as the predictive rule (see, e.g., Proust-Lima et al., 2014). However, this requires knowledge of the survival labels  $(T_i, \Delta_i)$ , which does not fit in our real-time prediction goal. Therefore, we define a new predictive marker as follows.

For any subject  $i$  and any time  $s_i$  elapsed since entry into the study, given longitudinal markers  $\mathcal{Y}_i(s_i^-)$  observed up to  $s_i$ , for any  $k \in \{1, \dots, K\}$ , we let

$$\widehat{\mathcal{R}}_{ik}(s_i) = \mathbb{P}_{\hat{\theta}}[g_i = k | T_i^* > s_i, \mathcal{Y}_i(s_i^-)].$$


Indeed, for any subject  $i$  who is event-free when  $s_i$  has elapsed, all we know about that subject is that its time to the event of interest  $T_i^*$  exceeds  $s_i$ . This is equivalent to considering this subject as a new subject for which  $T_i = s_i$ ,  $\Delta_i = 0$ , and  $Y_i = \mathcal{Y}_i(s_i^-)$ . The expression of  $\widehat{\mathcal{R}}_{ik}(s_i)$  can thus be rewritten as  $\widehat{\mathcal{R}}_{ik}(s_i) = \mathbb{P}_{\hat{\theta}}[g_i = k | T_i = s_i, \Delta_i = 0, Y_i = \mathcal{Y}_i(s_i^-)]$ , see Lemma 2 of the Supplementary Materials for more details.

We illustrate this real-time prediction setting in Figure 4, emphasizing that  $s_i$  should be thought of as the duration between the enrollment of individual  $i$  and the “present” time.

**4.1.2 Performance evaluation.** We want to compare the quality of our predictions to the true labels  $(T_i, \Delta_i)$ , to which we have access in these comparison experiments. Given a test set in which each individual’s trajectory is fully observed until the end of the study, we mimic the real-time prediction setting by randomly sampling the  $s_i$ .

We use the classical C-index (Harrell et al., 1996) as our performance metric. More precisely, we assume that we are in the case  $K = 2$  and the class  $g_i = 2$  represents the high-risk group of subjects (the class  $g_i = 1$  then representing a low-risk group). We denote by  $\widehat{\mathcal{R}}_i = \widehat{\mathcal{R}}_{i2}(s_i)$  the predictive marker that a subject  $i$  belongs to class  $g_i = 2$  when  $s_i$  has elapsed. Then,





**Figure 4:** Real-time prediction setting. In a practical application, we want to be able to make predictions at any “present” time while subjects have entered the study at different times. Therefore, some of them have a lot of recorded information while the others have a few.

we let  $\mathcal{C} = \mathbb{P}[\widehat{\mathcal{R}}_i > \widehat{\mathcal{R}}_j | T_i^* < T_j^*]$ , with  $i \neq j$  two random independent subjects (note that  $\mathcal{C}$  does not depend on  $i, j$  under the i.i.d. sample hypothesis).

In our case,  $T^*$  is subject to right censoring, so one would typically consider the modified  $\widetilde{\mathcal{C}}$  defined by  $\widetilde{\mathcal{C}} = \mathbb{P}[\widehat{\mathcal{R}}_i > \widehat{\mathcal{R}}_j | T_i < T_j, T_i < t^{\max}]$ , where  $t^{\max}$  corresponds to a fixed and predetermined follow-up period (Heagerty and Zheng, 2005). It has been shown by Uno et al. (2011) that a Kaplan-Meier estimator for the censoring distribution leads to a nonparametric and consistent estimator of  $\widetilde{\mathcal{C}}$ .

In Section C.4 of the Supplementary Materials, we give the complete procedure used to evaluate the performance of the models considered in our experiments.

#### 4.2 Competing models

We compare FLASH with the very classical and widely used LCMM (Proust-Lima et al., 2017) and `JMbayes` Rizopoulos (2016), which are extensions of JLCM and SREM that allow for multivariate longitudinal markers. We present them, along with their respective predictive markers, in Section B of the Supplementary Materials. Note that not many joint models allow for multivariate longitudinal markers, which limits our choice of competing methods.

### 5. Experimental results

To evaluate our method, we first perform in Subsection 5.1 a simulation study that illustrates our estimation procedure. We then turn to a comparison study on both simulated and medical examples in Subsection 5.2, and finally show in Subsection 5.3 that the biomarkers identified as significant by our model are consistent with current medical knowledge.


In all experiments, the features extracted by the `tsfresh` package (Christ et al., 2018) are used for association features  $\Psi_m$  in FLASH. This package extracts dozens of features from a time series such as absolute energy, kurtosis, or autocorrelation. Before running our extended EM algorithm with the set of features extracted by the `tsfresh` package, we use a screening phase procedure where we select the top ten association features by fitting the extracted feature of each candidate and the survival labels in individual Cox models and comparing their C-index scores. In addition, a recent line of work is to use the signature transform (Fermanian, 2021; Bleistein et al., 2024) to extract features from longitudinal markers. This transform encapsulates geometric information about multivariate time series. We provide additional results with the signature transform in Section C.7 of the Supplementary Materials, which show that our method is generic and performs well regardless of the feature extraction functions used.

We tune the regularization hyperparameters  $(\zeta_{1,k}, \zeta_{2,k})_{k \in \{1, \dots, K\}}$  with a grid search and a 10-fold cross-validation with the C-index metric. Note that we keep  $\zeta_{1,1} = \dots = \zeta_{1,K}$  and  $\zeta_{2,1} = \dots = \zeta_{2,K}$ . Extensive details on our experiments together with additional results on a high-dimensional dataset from NASA are given in Section C of the Supplementary Materials.


#### 5.1 Simulation study

To assess our estimation procedure, we simulate data as follows. First, we simulate the latent class membership indicator from the logistic regression model in (1). Based on this indicator, we divide the population into two groups: a high-risk group and a low-risk group. Within each group, we apply the classical survival simulation setting described by Bender et al. (2005). Next, we generate the longitudinal markers  $Y$  from the generalized linear mixed models in (3). Survival times are generated from their hazard functions given in (4). The model coefficients  $\xi_1$ ,  $\xi_2$ , and  $\gamma$  of models (1) and (4) are generated as sparse vectors, ensuring that only a subset of the corresponding features are active (i.e., the coefficients are non-zero). Extensive details of these simulations are provided in Section C.2 of the Supplementary Materials. Figure 5 shows some examples of simulated longitudinal markers and Kaplan-Meier survival curves.

To illustrate our regularization strategy, we show in Figure 6 the time-independent parameter  $\xi$  and the joint association



**Figure 5:** Simulated cohort of  $n = 500$  samples for  $K = 2$  groups (high-risk group in red curves and low-risk group in blue curves). Top figures: trajectories of first three longitudinal markers of two individuals randomly selected in each group. Bottom figure: Kaplan-Meier survival curves for each group.



**Figure 6:** Simulations results. (a): the support of both the true coefficient  $\xi$  in green and its estimated version  $\hat{\xi}$  in red. (b) and (c): in red the support of the estimated coefficient  $\hat{\gamma}_k$  for  $k \in \{1, 2\}$ , the dashed pink lines separate the features corresponding to each longitudinal marker  $\ell$ , and active longitudinal markers are represented by a green area.

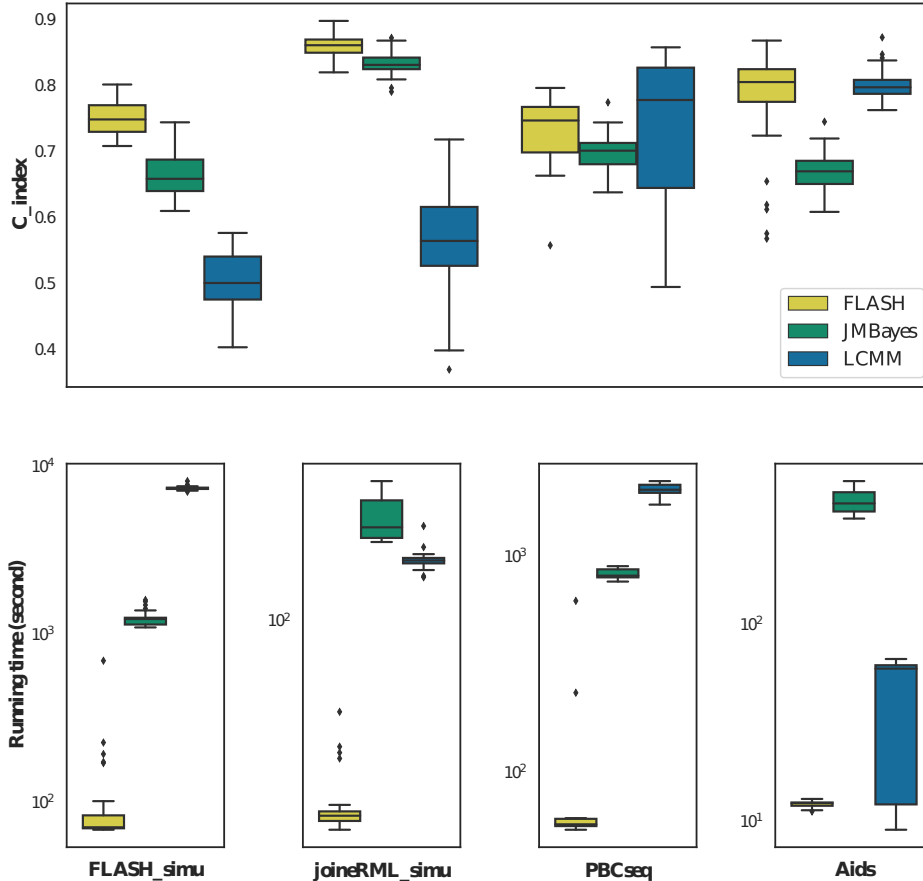
parameters  $(\gamma_k)_{k \in \{1, 2\}}$  and their estimation obtained after running our learning procedure. We see in the sub-figure (a) that the support of  $\xi$  is fully recovered thanks to the elastic-net penalty. Additionally, sub-figures (b) and (c) demonstrate the effect of the sparse group lasso, showing that only the coefficients corresponding to active longitudinal features (represented by the green area) are non-zero, while all coefficients for inactive longitudinal features are zero.

### 5.2 Comparison study

We compare FLASH with `JMbayes` and `LCMM` on two simulated and two real-world datasets and use the C-index metric presented in Section 4. The first simulated dataset is the one from the previous subsection, and the second one is from the R package (Hickey et al., 2018) `joineRML`. We give detailed description of this second simulated dataset and the two real-world datasets (*PBCseq* and *Aids*) in Section C.3 of the Supplementary Materials. A summary of the datasets is given in Table 1.

Table 1: Datasets characteristics: the number of samples  $n$ , the number of longitudinal features  $L$ , number of time-independent features  $p$ , and the overall number of parameters in FLASH model  $P$ . The names `FLASH_simu` and `joineRML_simu` correspond to the datasets simulated from the simulation study in Section 5.1 and the `joineRML` package respectively.

Dataset	$n$	$L$	$p$	$P$
<code>FLASH_simu</code>	500	5	10	224
<code>joineRML_simu</code>	250	2	2	204
<i>PBCseq</i>	304	7	3	251
<i>Aids</i>	467	1	4	147
<i>Sepsis</i>	654	4	21	255



**Figure 7:** C-index (**top** figure) and runtime (**bottom** figures) comparison on the four datasets considered. The box plots of C-index and runtime are obtained with 50 independent experiments.

We can see in Figure 7 that FLASH outperforms its competitors in terms of both C-index and running times on all datasets. The good performance of FLASH in terms of running times can be explained by the fact that it does not need to perform computationally intensive Monte Carlo techniques like `JMbayes`, while it is easier to satisfy the convergence criterion of our EM algorithm than that of `LCMM`.

### 5.3 Biological interpretation of FLASH results

Considering the growing emphasis on model interpretability and the fact that the new regulations in the European Union and the United States now require that a model be interpretable and understandable to be certified as a medical device (Geller, 2023; Panigutti et al., 2023), we conclude this section with interpretations of the results of FLASH on the PBCseq and Sepsis datasets. The Sepsis dataset describes the sepsis diagnosis of patients, where, after a pre-processing step, 4 multivariate longitudinal features and 21 time-independent features are available for each patient. The coefficients estimated by FLASH, and in particular their sparsity, provide us information on which marker is involved in the diagnosis, see Section C.5 of the Supplementary Materials for all numerical values. Note that we did not include the Sepsis dataset in Section 5.2 because both LCMM and JMBayes take a very long time to converge, so that we could not do the Monte Carlo comparison. This highlights the fact that they do not scale to high-dimensional settings. However, on one experimental trial, LCMM and JMBayes only give C-index scores of 0.51 and 0.56 while FLASH has a score of 0.74.

For the PBCseq dataset, alkaline phosphatase appears to be the most important variable, followed by prothrombin and albumin. Alkanine phosphatase is already recognized to be an important variable to monitor: it is already known that phosphatase alkaline at 6-month predicts non-responders and survival (Perez et al., 2023), it is recognized that treatment target should be normalization of alkaline phosphatase (Perez et al., 2020), and prognosis is improved for patients taking drugs lowering alkaline phosphatase as ursodeoxycholic acid (Kuiper et al., 2009). Concerning the Sepsis dataset, respiratory rate is the most important variable. Systolic blood pressure appears to be the most important longitudinal feature, followed by oxygen saturation. Of note, the two most widely used prognostic criteria in sepsis, i.e. the qSOFA and the SIRS criteria, which contain respectively 3 and 4 variables, both include among these variables respiratory rate (Raith et al., 2017).

## 6. Discussion

In this paper, a generalized joint model for high-dimensional multivariate longitudinal data and censored durations (FLASH) has been introduced, with an efficient estimation methodology based on an extension of the EM algorithm. This algorithm allows the use of regularization strategies in order to perform feature selection and results in an interpretable model scalable to high-dimensional longitudinal markers. We evaluated the performance of the estimation procedure on an extensive Monte Carlo simulation study. It showed that our method successfully recovered the most significant features. The proposed methodology has then been applied on four different datasets. On these datasets, FLASH outperforms competing methods, both in terms of C-index and runtimes, in a so-called “real-time” prediction setting. In addition, we show on experiments of medical datasets that our model automatically identifies the most important longitudinal markers and time-independent features, allowing important interpretations on the application at hand. Potential future work consists of extending the implementation to generalize our EM algorithm to support count or binary longitudinal features, and to relax the assumptions on latent class membership to allow classes to change with time, using for example a Markov structure.

## ACKNOWLEDGEMENTS

*Conflict of Interest:* None declared. For this work, Agathe Guilloux and Van Tuan Nguyen benefited from the support of the Agence Nationale pour la Recherche under the France 2030 program with the reference ANR-22-PESN-0016. This work was also supported by the French National Cancer Institut (INCa) [grant number 2016-1-PL SHS-03-1]. The authors thank Linus Bleistein and Massil Hihat for fruitful discussions.

## SUPPLEMENTARY MATERIALS

Supplementary Materials, which includes the sections referenced in the main paper, is available with this paper at the Biometrics website on Wiley Online Library.

## REFERENCES

- Aldrich, J. H. and Nelson, F. D. (1984). *Linear probability, logit, and probit models*. Number 45. Sage.
- Andersen, P. K. and Gill, R. D. (1982). Cox’s regression model for counting processes: a large sample study. *The annals of statistics* pages 1100–1120.
- Andrew, G. and Gao, J. (2007). Scalable training of  $l_1$ -regularized log-linear models. In *International Conference on Machine Learning*, pages 33–40. ACM.
- Andrinopoulou, E.-R., Nasserinejad, K., Szczesniak, R., and Rizopoulos, D. (2020). Integrating latent classes in the bayesian shared parameter joint model of longitudinal and survival outcomes. *Statistical methods in medical research* **29**, 3294–3307.
- Andrinopoulou, E.-R. and Rizopoulos, D. (2016). Bayesian shrinkage approach for a joint model of longitudinal and survival outcomes assuming different association structures. *Statistics in medicine* **35**, 4813–4823.

- Austin, P. C. (2013). Correction: 'generating survival times to simulate cox proportional hazards models with time-varying covariates'. *Statistics in Medicine* **32**, 1078–1078.
- Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al. (2012). Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning* **4**, 1–106.
- Bach, F. R. (2008). Bolasso: model consistent lasso estimation through the bootstrap. In *Proceedings of the 25th international conference on Machine learning*, pages 33–40.
- Bartolucci, F. and Farcomeni, A. (2015). A discrete time event-history approach to informative drop-out in mixed latent markov models with covariates. *Biometrics* **71**, 80–89.
- Bartolucci, F. and Farcomeni, A. (2019). A shared-parameter continuous-time hidden markov and survival model for longitudinal data with informative dropout. *Statistics in medicine* **38**, 1056–1073.
- Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences* **2**, 183–202.
- Bender, R., Augustin, T., and Blettner, M. (2005). Generating survival times to simulate cox proportional hazards models. *Statistics in medicine* **24**, 1713–1723.
- Bleistein, L., Nguyen, V.-T., Fermanian, A., and Guillaoux, A. (2024). Dynamical survival analysis with controlled latent states. *41th International Conference on Machine Learning and arXiv preprint arXiv:2401.17077*.
- Böhning, D. (1992). Multinomial logistic regression algorithm. *Annals of the institute of Statistical Mathematics* **44**, 197–200.
- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press, New York.
- Breslow, N. E. (1972). Contribution to the discussion of the paper by dr cox. *Journal of the Royal Statistical Society, Series B* **34**, 216–217.
- Bussy, S., Guillaoux, A., Gaïffas, S., and Jannot, A.-S. (2019). C-mix: A high-dimensional mixture model for censored durations, with applications to genetic data. *Statistical methods in medical research* **28**, 1523–1539.
- Chi, Y.-Y. and Ibrahim, J. G. (2006). Joint models for multivariate longitudinal and multivariate survival data. *Biometrics* **62**, 432–445.
- Christ, M., Braun, N., Neuffer, J., and Kempa-Liehr, A. W. (2018). Time series feature extraction on basis of scalable hypothesis tests (tsfresh—a python package). *Neurocomputing* **307**, 72–77.
- Chzhen, E., Hebiri, M., and Salmon, J. (2019). On lasso refitting strategies.
- Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)* **34**, 187–220.
- Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* **39**, 1–38.
- Devaux, A., Genuer, R., Peres, K., and Proust-Lima, C. (2022). Individual dynamic prediction of clinical endpoint from large dimensional longitudinal biomarker history: a landmark approach. *BMC Medical Research Methodology* **22**, 1–14.
- Fabian Pedregosa, Geoffrey Negiar, G. D. (2020). copt: composite optimization in python.
- Fabio, L. C., Paula, G. A., and de Castro, M. (2012). A poisson mixed model with nonnormal random effect distribution. *Computational Statistics & Data Analysis* **56**, 1499–1510.
- Fermanian, A. (2021). Embedding and learning with signatures. *Computational Statistics & Data Analysis* **157**, 107148.
- Fitzmaurice, G. M., Laird, N. M., and Ware, J. H. (2012). *Applied longitudinal analysis*. John Wiley & Sons, New Jersey.
- Geller, J. (2023). Food and drug administration published final guidance on clinical decision support software. *Journal of Clinical Engineering* **48**, 3–7.
- Gompertz, B. (1825). Xxiv. on the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. in a letter to francis baily, esq. frs &c. *Philosophical transactions of the Royal Society of London* pages 513–583.
- Harrell, F. E., Lee, K. L., and Mark, D. B. (1996). Tutorial in biostatistics multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in medicine* **15**, 361–387.
- Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.
- Hatfield, L. A., Boye, M. E., and Carlin, B. P. (2011). Joint modeling of multiple longitudinal patient-reported outcomes and survival. *Journal of Biopharmaceutical Statistics* **21**, 971–991.
- Heagerty, P. J. and Zheng, Y. (2005). Survival model predictive accuracy and roc curves. *Biometrics* **61**, 92–105.
- Hickey, G. L., Philipson, P., Jorgensen, A., and Kolamunnage-Dona, R. (2016). Joint modelling of time-to-event and multivariate longitudinal outcomes: recent developments and issues. *BMC medical research methodology* **16**, 117.
- Hickey, G. L., Philipson, P., Jorgensen, A., and Kolamunnage-Dona, R. (2018). joinerml: a joint model and software package for time-to-event and multivariate longitudinal outcomes. *BMC medical research methodology* **18**, 1–14.
- Jaffa, M. A., Gebregziabher, M., and Jaffa, A. A. (2014). A joint modeling approach for right censored high dimensional multivariate longitudinal data. *Journal of biometrics & biostatistics* **5**.
- Kang, K. and Song, X. (2022). Consistent estimation of a joint model for multivariate longitudinal and survival data with latent variables. *Journal of Multivariate Analysis* **187**, 104827.
- Klein, J. P. (1992). Semiparametric estimation of random effects using the cox model based on the em algorithm. *Biometrics* pages 795–806.

- Klein, J. P. and Moeschberger, M. L. (2005). *Survival analysis: techniques for censored and truncated data*. Springer Science & Business Media, New York.
- Kuiper, E. M., Hansen, B. E., de Vries, R. A., den Ouden-Muller, J. W., Van Ditzhuijsen, T. J., Haagsma, E. B., Houben, M. H., Wittteman, B. J., van Erpecum, K. J., van Buuren, H. R., et al. (2009). Improved prognosis of patients with primary biliary cirrhosis that have a biochemical response to ursodeoxycholic acid. *Gastroenterology* **136**, 1281–1287.
- Lin, H., McCulloch, C. E., and Mayne, S. T. (2002). Maximum likelihood estimation in the joint analysis of time-to-event and multiple longitudinal variables. *Statistics in Medicine* **21**, 2369–2382.
- Lin, H., Turnbull, B. W., McCulloch, C. E., and Slate, E. H. (2002). Latent class models for joint analysis of longitudinal biomarker and event process data: application to longitudinal prostate-specific antigen readings and prostate cancer. *Journal of the American Statistical Association* **97**, 53–65.
- Masoudnia, S. and Ebrahimpour, R. (2014). Mixture of experts: a literature survey. *Artificial Intelligence Review* **42**, 275–293.
- McLachlan, G. J. and Krishnan, T. (2007). *The EM algorithm and extensions*. John Wiley & Sons.
- Molenberghs, G., Verbeke, G., Demétrio, C. G., and Vieira, A. M. (2010). A family of generalized linear models for repeated measures with normal and conjugate random effects.
- Moreau, J. J. (1962). Fonctions convexes duales et points proximaux dans un espace hilbertien. *Comptes rendus hebdomadaires des séances de l'Académie des sciences* **255**, 2897–2899.
- Mukherjee, B. N. and Maiti, S. S. (1988). On some properties of positive definite toeplitz matrices and their possible applications. *Linear algebra and its applications* **102**, 211–240.
- Murray, J. and Philipson, P. (2022). A fast approximate em algorithm for joint models of survival and multivariate longitudinal data. *Computational Statistics & Data Analysis* **170**, 107438.
- Panigutti, C., Hamon, R., Hupont, I., Fernandez Llorca, D., Fano Yela, D., Junklewitz, H., Scalzo, S., Mazzini, G., Sanchez, I., Soler Garrido, J., et al. (2023). The role of explainable ai in the context of the ai act. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1139–1150.
- Perez, C. F. M., Harms, M. H., Lindor, K. D., Van Buuren, H. R., Hirschfield, G. M., Corpechot, C., Van Der Meer, A. J., Feld, J. J., Gulamhusein, A., Lammers, W. J., et al. (2020). Goals of treatment for improved survival in primary biliary cholangitis: treatment target should be bilirubin within the normal range and normalization of alkaline phosphatase. *Official journal of the American College of Gastroenterology— ACG* **115**, 1066–1074.
- Perez, C. F. M., Ioannou, S., Hassanally, I., Trivedi, P. J., Corpechot, C., van der Meer, A. J., Lammers, W. J., Battezzati, P. M., Lindor, K. D., Nevens, F., et al. (2023). Optimizing therapy in primary biliary cholangitis: Alkaline phosphatase at six months identifies one-year non-responders and predicts survival. *Liver international: official journal of the International Association for the Study of the Liver* **43**, 1497–1506.
- Proust-Lima, C., Philipps, V., and Liquet, B. (2017). Estimation of extended mixed models using latent classes and latent processes: The r package lcmm. *Journal of Statistical Software, Articles* **78**, 1–56.
- Proust-Lima, C., Séne, M., Taylor, J. M., and Jacqmin-Gadda, H. (2014). Joint latent class models for longitudinal and time-to-event data: A review. *Statistical methods in medical research* **23**, 74–90.
- Raith, E. P., Udy, A. A., Bailey, M., McGloughlin, S., MacIsaac, C., Bellomo, R., Pilcher, D. V., et al. (2017). Prognostic accuracy of the sofa score, sirs criteria, and qsofa score for in-hospital mortality among adults with suspected infection admitted to the intensive care unit. *Jama* **317**, 290–300.
- Rizopoulos, D. (2016). The r package jmbayes for fitting joint models for longitudinal and time-to-event data using mcmc. *Journal of Statistical Software, Articles* **72**, 1–46.
- Rizopoulos, D. and Ghosh, P. (2011). A bayesian semiparametric multivariate joint model for multiple longitudinal outcomes and a time-to-event. *Statistics in medicine* **30**, 1366–1380.
- Rustand, D., Van Niekerk, J., Krainski, E. T., Rue, H., and Proust-Lima, C. (2024). Fast and flexible inference for joint models of multivariate longitudinal and survival data using integrated nested laplace approximations. *Biostatistics* **25**, 429–448.
- Simon, N., Friedman, J., Hastie, T., and Tibshirani, R. (2013). A sparse-group lasso. *Journal of Computational and Graphical Statistics* **22**, 231–245.
- Tsiatis, A. A. and Davidian, M. (2004). Joint modeling of longitudinal and time-to-event data: an overview. *Statistica Sinica* pages 809–834.
- Uno, H., Cai, T., Pencina, M. J., D’Agostino, R. B., and Wei, L. (2011). On the c-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in medicine* **30**, 1105–1117.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods* **17**, 261–272.
- Wang, P., Shen, W., and Boye, M. E. (2012). Joint modeling of longitudinal outcomes and survival using latent growth modeling approach in a mesothelioma trial. *Health Services and Outcomes Research Methodology* **12**, 182–199.
- Wulfsohn, M. S. and Tsiatis, A. A. (1997). A joint model for survival and longitudinal data measured with error. *Biometrics* pages 330–339.
- Yu, M., Law, N. J., Taylor, J. M., and Sandler, H. M. (2004). Joint longitudinal-survival-cure models and their application to prostate cancer. *Statistica Sinica* pages 835–862.

- Yuan, L., Liu, J., and Ye, J. (2011). Efficient methods for overlapping group lasso. *Advances in neural information processing systems* **24**, 352–360.
- Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. (1997). Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)* **23**, 550–560.
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**, 301–320.

## Supplementary Material

To help the reader, Table 2 provides a list of notations used in the the paper.

Table 2: Summary of notation

Notation	Definition
$X_i$	Time-independent feature
$Y_i$	Longitudinal markers
$T_i$	Survival time
$\Delta_i$	Censoring indication
$g_i$	Latent class membership variable
$b_i$	Random effects of the longitudinal markers
$n$	Number of training samples
$p$	Number of time-independent features
$L$	Number of longitudinal markers
$K$	Number of latent classes
$D$	Variance-covariance matrix of the $b_i$
$U_i$	Fixed-effect design matrix
$V_i$	Random-effect design matrix
$\psi_i$	Association features
$\xi$	Time-independent parameters
$\beta$	Fixed-effect parameter
$\gamma$	Joint association parameter
$\lambda_0$	Baseline hazard function
$I_m$	Identity matrix of $\mathbb{R}^{m \times m}$
$\mathbf{1}_m$	Vector of $\mathbb{R}^m$ having all coordinates equal to one
$\mathbf{0}_m$	Vector of $\mathbb{R}^m$ having all coordinates equal to zero
$\ \cdot\ _q$	The usual $\ell_q$ -quasi norm, $q > 0$

### A. DETAILS ON THE EXTENDED EM ALGORITHM

We detail in this section our learning methodology. First, recall that the penalized negative log-likelihood is defined by

$$\mathcal{L}_n^{\text{pen}}(\theta) = \mathcal{L}_n(\theta) + \sum_{k=1}^K \zeta_{1,k} \Omega_1(\xi_k) + \sum_{k=1}^K \zeta_{2,k} \Omega_2(\gamma_k), \quad (1)$$

where

$$\Omega_1(\xi_k) = (1 - \eta) \|\xi_k\|_1 + \frac{\eta}{2} \|\xi_k\|_2^2 \quad \text{and} \quad \Omega_2(\gamma_k) = (1 - \tilde{\eta}) \|\gamma_k\|_1 + \tilde{\eta} \sum_{\ell=1}^L \|\gamma_k^\ell\|_2,$$

where  $(\eta, \tilde{\eta}) \in [0, 1]^2$  are fixed (depending on the level of sparsity expected),  $\gamma_k^\ell = (\gamma_{k,1}^\ell, \dots, \gamma_{k,M}^\ell)^\top \in \mathbb{R}^M$  is the subset of  $\gamma_k$  corresponding to the longitudinal marker  $\ell$ ,  $\|\cdot\|_1$  (resp.  $\|\cdot\|_2$ ) denotes the usual  $\ell_1$  (resp.  $\ell_2$ ) norm. In all our experiments, we take  $\eta = 0.1$  and  $\tilde{\eta} = 0.9$ .

The goal is to minimize this objective function by the EM algorithm. This is done in two steps: compute the expectation of the negative complete log-likelihood with respect to the unobserved quantities (the random effects  $b_i$  and the latent classes  $g_i$ ), then minimize the obtained quantity with respect to all parameters of the model in  $\theta$ . For simplicity, we won't compute all terms of the expectation in the E-step but only the quantities used in the M-step. Moreover, we perform minimization with respect to  $\theta$  in several steps, minimizing with respect to block of parameters separately to obtain tractable updates.



## A.1 E-step

Recall that, under our assumptions, the negative complete log-likelihood writes

$$\mathcal{L}_n^{\text{comp}}(\theta) = -n^{-1} \sum_{i=1}^n \left( \log f_\theta(b_i) + \sum_{k=1}^K \mathbb{1}_{\{g_i=k\}} (\log \mathbb{P}_\theta(g_i = k) + \log f_\theta(Y_i | b_i, g_i = k)) \right) \quad (2)$$

$$+ \log f_\theta(T_i, \Delta_i | Y_i, g_i = k). \quad (3)$$

Let us introduce a few matrix notations. Concatenating all longitudinal markers and all observation times, the mean of the vector  $Y_i | b_i, g_i = k$  (defined in (3) in the main paper) can be rewritten  $M_{ik} = U_i \beta_k + V_i b_i$ , where we introduce the design matrices

$$U_i = \begin{bmatrix} U_i^1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & U_i^L \end{bmatrix} \in \mathbb{R}^{n_i \times q} \quad \text{and} \quad V_i = \begin{bmatrix} V_i^1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & V_i^L \end{bmatrix} \in \mathbb{R}^{n_i \times r}$$

and for all  $\ell \in \{1, \dots, L\}$ , one writes

$$U_i^\ell = \begin{bmatrix} u^\ell(t_{i1}^\ell)^\top \\ \vdots \\ u^\ell(t_{in_i^\ell}^\ell)^\top \end{bmatrix} \in \mathbb{R}^{n_i^\ell \times q_\ell} \quad \text{and} \quad V_i^\ell = \begin{bmatrix} v^\ell(t_{i1}^\ell)^\top \\ \vdots \\ v^\ell(t_{in_i^\ell}^\ell)^\top \end{bmatrix} \in \mathbb{R}^{n_i^\ell \times r_\ell}.$$

Under all assumptions of Section 2 in the main paper, we can then write explicitly the different terms. The random effects simply follow a Gaussian distribution, which yields

$$\log f_\theta(b_i) = -\frac{1}{2} (r \log 2\pi + \log \det(D) + b_i^\top D^{-1} b_i).$$

Then, the conditional density of the longitudinal features (in the Gaussian case) writes

$$\log f_\theta(Y_i | b_i, g_i = k) = -\frac{1}{2} (n_i \log 2\pi + \log \det(\Sigma_i) + (Y_i - U_i \beta_k - V_i b_i)^\top \Sigma_i^{-1} (Y_i - U_i \beta_k - V_i b_i)).$$

Finally, the survival terms in the complete likelihood write

$$\log f_\theta(T_i, \Delta_i | Y_i, g_i = k) = \Delta_i (\log \lambda_0(T_i) + \psi_i(T_i)^\top \gamma_k) - \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k) \mathbb{1}_{\{\tau_j \leq T_i\}}. \quad (4)$$

We can therefore decompose the expected negative complete log-likelihood as

$$\begin{aligned} \mathcal{Q}_n(\theta, \theta^{(w)}) &= \mathbb{E}_{\theta^{(w)}} [\mathcal{L}_n^{\text{comp}}(\theta) | \mathcal{D}_n] \\ &= -n^{-1} \sum_{i=1}^n \left( A_i^1(D) + \sum_{k=1}^K \mathbb{P}_{\theta^{(w)}}(g_i = k | \mathcal{D}_n) (A_{i,k}^2(\xi) + A_i^3(\beta_k, \Sigma_i) + A_i^4(\gamma_k, \lambda_0)) \right) \\ &\quad + \text{constants}, \end{aligned}$$

where, for any  $b \in \mathbb{R}^r$ ,  $D \in \mathbb{R}^{r \times r}$ ,  $\beta \in \mathbb{R}^q$ ,  $\Sigma \in \mathbb{R}^{n_i}$ , we define

$$\begin{aligned} A_i^1(D) &= \frac{1}{2} \mathbb{E}_{\theta^{(w)}} [\log \det(D) + b_i^\top D^{-1} b_i | \mathcal{D}_n], \\ A_{i,k}^2(\xi) &= \log \left( \frac{e^{X_i^\top \xi_k}}{\sum_{j=1}^K e^{X_i^\top \xi_j}} \right), \\ A_i^3(\beta, \Sigma) &= \frac{1}{2} \mathbb{E}_{\theta^{(w)}} [\log \det(\Sigma) + (Y_i - U_i \beta - V_i b_i)^\top \Sigma^{-1} (Y_i - U_i \beta - V_i b_i) | \mathcal{D}_n], \\ A_i^4(\gamma, \lambda_0) &= \Delta_i (\log \lambda_0(T_i) + \psi_i(T_i)^\top \gamma) - \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma) \mathbb{1}_{\{\tau_j \leq T_i\}}. \end{aligned}$$

As explained before, we do not compute the expectation of the negative complete likelihood but only of the quantities needed in the M-step. First, the following technical lemmas will prove useful in the E-step and makes use of conjugate properties of Gaussian distributions.

LEMMA 1: For any  $i \in \{1, \dots, n\}$ ,  $k \in \{1, \dots, K\}$ ,

$$Y_i | g_i = k \sim \mathcal{N}(U_i \beta_k, V_i D V_i^\top + \Sigma_i) \quad \text{and} \quad b_i | Y_i, g_i = k \sim \mathcal{N}(O_{i,k}, W_i),$$

where

$$O_{i,k} = W_i V_i^T \Sigma_i^{-1} (Y_i - U_i \beta_k) \quad \text{and} \quad W_i = (V_i^T \Sigma_i^{-1} V_i + D^{-1})^{-1}. \quad (5)$$

*Proof.* From (3) in the main paper we know that  $Y_i | b_i, g_i = k \sim \mathcal{N}(M_{ik}, \Sigma_i)$  and  $b_i \sim \mathcal{N}(0, D)$ , which gives

$$Y_i | g_i = k \sim \mathcal{N}(U_i \beta_k, V_i D V_i^T + \Sigma_i).$$

Moreover, by Bayes's rule, the distribution of  $b_i | Y_i, g_i = k$  can be written

$$\begin{aligned} f_\theta(b_i | Y_i, g_i = k) &\propto f_\theta(Y_i | b_i, g_i = k) f_\theta(b_i | g_i = k) \\ &\propto \exp((Y_i - U_i \beta_k - V_i b_i)^T \Sigma_i^{-1} (Y_i - U_i \beta_k - V_i b_i) + b_i^T D^{-1} b_i) \\ &\propto \exp((b_i - O_{i,k})^T W_i^{-1} (b_i - O_{i,k})), \end{aligned}$$

where  $O_{i,k} = W_i V_i^T \Sigma_i^{-1} (Y_i - U_i \beta_k)$  and  $W_i = (V_i^T \Sigma_i^{-1} V_i + D^{-1})^{-1}$ . We then have

$$b_i | Y_i, g_i = k \sim \mathcal{N}(O_{i,k}, W_i).$$

The following lemma gives the three expectations that appear in this E-step.

LEMMA 2: For any  $i \in \{1, \dots, n\}$ ,  $k \in \{1, \dots, K\}$ ,  $\theta \in \mathbb{R}^P$ , the three following integrals are closed-form and write

$$\mathbb{E}_\theta[b_i | \mathcal{D}_n] = \frac{\sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(Y_i | g_i = j) O_{i,j}}{\sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(Y_i | g_i = j)}, \quad (6)$$

$$\mathbb{E}_\theta[b_i b_i^\top | \mathcal{D}_n] = \frac{\sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(Y_i | g_i = j) (W_i + O_{i,j} O_{i,j}^\top)}{\sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(Y_i | g_i = j)}, \quad (7)$$

and

$$\tilde{\pi}_{ik}^\theta = \mathbb{P}_\theta(g_i = k | \mathcal{D}_n) = \frac{\mathbb{P}_\theta(g_i = k) f_\theta(T_i, \Delta_i | Y_i, g_i = k) f_\theta(Y_i | g_i = k)}{\sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(Y_i | g_i = j)}, \quad (8)$$

where  $f_\theta(Y_i | g_i = j)$  is the density of the multivariate Gaussian distribution of Lemma 1.

*Proof.* By Assumption 2 in the main paper, the distribution of  $b_i$  given the observed data  $\mathcal{D}_n$ , for any  $\theta \in \mathbb{R}^P$ , writes

$$\begin{aligned} f_\theta(b_i | \mathcal{D}_n) &= \frac{f_\theta(b_i, T_i, \Delta_i, Y_i)}{f_\theta(T_i, \Delta_i, Y_i)} \\ &= \frac{1}{f_\theta(T_i, \Delta_i, Y_i)} \sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(b_i, T_i, \Delta_i, Y_i | g_i = j) \\ &= \frac{1}{f_\theta(T_i, \Delta_i, Y_i)} \sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(Y_i | g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) f_\theta(b_i | Y_i, g_i = j). \end{aligned}$$

Similarly, we have

$$f_\theta(T_i, \Delta_i, Y_i) = \sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(Y_i | g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j).$$

This gives that, for any function  $\mu$ ,

$$\begin{aligned} \mathbb{E}_\theta[\mu(b_i) | \mathcal{D}_n] &= \int_{\mathbb{R}^r} \mu(b_i) f_\theta(b_i | \mathcal{D}_n) db_i \\ &= \frac{1}{f_\theta(T_i, \Delta_i, Y_i)} \sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(Y_i | g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) \int_{\mathbb{R}^r} \mu(b_i) f_\theta(b_i | Y_i, g_i = j) db_i \\ &= \frac{1}{f_\theta(T_i, \Delta_i, Y_i)} \sum_{j=1}^K \mathbb{P}_\theta(g_i = j) f_\theta(Y_i | g_i = j) f_\theta(T_i, \Delta_i | Y_i, g_i = j) \mathbb{E}_\theta[\mu(b_i) | Y_i, g_i = j] \end{aligned}$$

Hence, for the functions  $\mu(b_i) = b_i$  and  $\mu(b_i) = b_i b_i^\top$ , we obtain the result by applying Lemma 1, which gives

$$\mathbb{E}_\theta[b_i | Y_i, g_i = j] = O_{i,j}, \quad \text{and} \quad \mathbb{E}_\theta[b_i b_i^\top | Y_i, g_i = j] = W_i + O_{i,j} O_{i,j}^\top.$$

In the same manner, we can see that for any  $k \in \{1, \dots, K\}$ ,

$$\begin{aligned} \mathbb{E}_\theta[\mathbb{1}_{\{g_i=k\}}|T_i, \Delta_i, Y_i] &= \mathbb{P}_\theta(g_i = k|T_i, \Delta_i, Y_i) \\ &= \frac{1}{f_\theta(T_i, \Delta_i, Y_i)} \mathbb{P}_\theta(g_i = k) f_\theta(T_i, \Delta_i|Y_i, g_i = k) f_\theta(Y_i|g_i = k), \end{aligned}$$

which concludes the proof.

*Extensions of the longitudinal model.* To support other types of longitudinal features (e.g., count, binary, etc) where the conditional longitudinal features  $Y_i|b_i, g_i = k$  are not Gaussian, the simple solution is to choose distributions such that the random effect  $b_i$  and  $Y_i|b_i, g_i = k$  have conjugate distributions. There are several options, for example using the beta-binomial distribution for binary features or negative-binomial for count features. In this case, the expectation  $\mathbb{E}_\theta[\mu(b_i)|Y_i, g_i = j]$  and the density  $f_\theta(b_i|Y_i, g_i = j)$  can be computed in closed forms (Molenberghs et al., 2010) and the extension is trivial. If the distributions are not conjugate, some numerical integration methods can be used to compute  $\mathbb{E}_\theta[\mu(b_i)|Y_i, g_i = j]$  and  $f_\theta(b_i|Y_i, g_i = j)$  (see, e.g., Fabio et al., 2012).

## A.2 $M$ -step: closed-form updates

Now, we assume that we are at step  $w + 1$  of the algorithm, meaning that we have a current value  $\theta^{(w)}$  of the parameters and we update it to get the new parameters  $\theta^{(w+1)}$  by solving

$$\theta^{(w+1)} \in \underset{\theta \in \mathbb{R}^P}{\operatorname{argmin}} \mathcal{Q}_n^{\text{pen}}(\theta, \theta^{(w)}). \quad (9)$$

Note that we update for the coordinates of  $\theta^{(w)}$  in order, which is  $(D^{(w+1)}, (\xi_k^{(w+1)})_{k \in \{1, \dots, K\}}, (\beta_k^{(w+1)})_{k \in \{1, \dots, K\}}, (\gamma_k^{(w+1)})_{k \in \{1, \dots, K\}}, \lambda_0^{(w+1)}, \phi^{(w+1)})$ . Then, following this order, the update of later coordinates use the latest update of the previous ones. The update of several coordinates of  $\theta^{(w+1)}$  can be obtained in closed-form.

LEMMA 3: *At step  $w + 1$  of the EM algorithm, the update of  $D$  is*

$$D^{(w+1)} = n^{-1} \sum_{i=1}^n \mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i],$$

where  $\mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i]$  is given by (7).

*Proof.* The update for  $D^{(w)}$  requires to solve the following minimization problem

$$D^{(w+1)} \in \underset{D \in \mathbb{R}^{q \times q}}{\operatorname{argmin}} -n^{-1} \sum_{i=1}^n A_i^1(D), \quad (10)$$

We have for any  $i \in \{1, \dots, n\}$ ,

$$A_i^1(D) = \mathbb{E}_{\theta^{(w)}}[\log \det(D) + b_i^\top D^{-1} b_i] = \log \det(D) + \int_{\mathbb{R}^r} b_i^\top D^{-1} b_i f_{\theta^{(w)}}(b_i | T_i, \Delta_i, Y_i) db_i.$$

The gradient of  $A_i^1$  is here given by

$$\begin{aligned} \frac{\partial A_i^1(D)}{\partial D} &= D^{-\top} - \int_{\mathbb{R}^r} D^{-\top} b_i b_i^\top D^{-\top} f_{\theta^{(w)}}(b_i | T_i, \Delta_i, Y_i) db_i \\ &= D^{-\top} - D^{-\top} \mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i] D^{-\top}, \end{aligned}$$

where  $D^{-\top}$  is the transpose of matrix  $D^{-1}$ . The proof is completed by cancelling the gradient, that is

$$-n^{-1} \sum_{i=1}^n \frac{\partial A_i^1(D)}{\partial D} = 0 \quad \Leftrightarrow \quad -D^{-\top} + n^{-1} \sum_{i=1}^n D^{-\top} \mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i] D^{-\top} = 0$$

LEMMA 4: *At step  $w + 1$  of the EM algorithm, the update of  $\beta_k$  is*

$$\beta_k^{(w+1)} = \left( \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} U_i^\top U_i \right)^{-1} \left( \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} U_i^\top (Y_i - V_i \mathbb{E}_{\theta^{(w)}}[b_i | T_i, \Delta_i, Y_i]) \right), \quad (11)$$

where  $\mathbb{E}_{\theta^{(w)}}[b_i | T_i, \Delta_i, Y_i]$  and  $\tilde{\pi}_{ik}^{\theta^{(w)}}$  are given respectively by (6) and (8).

*Proof.* The update for  $\beta_k^{(w)}$  requires to solve the following minimization problem

$$\beta_k^{(w+1)} \in \operatorname{argmin}_{\beta \in \mathbb{R}^q} -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} A_i^3(\beta, \Sigma_i^{(w)}). \quad (12)$$

We have

$$\begin{aligned} A_i^3(\beta, \Sigma_i^{(w)}) &= \frac{1}{2} \mathbb{E}_{\theta^{(w)}} [\log \det(\Sigma_i^{(w)}) + (Y_i - U_i \beta - V_i b_i)^\top \Sigma_i^{(w)-1} (Y_i - U_i \beta - V_i b_i) | \mathcal{D}_n] \\ &= \frac{1}{2} \log \det(\Sigma_i^{(w)}) + \frac{1}{2} \mathbb{E}_{\theta^{(w)}} [(Y_i - U_i \beta - V_i b_i)^\top \Sigma_i^{(w)-1} (Y_i - U_i \beta - V_i b_i) | \mathcal{D}_n] \\ &= \mathbb{E}_{\theta^{(w)}} [- (Y_i - V_i b_i)^\top \Sigma_i^{(w)-1} U_i \beta + \frac{1}{2} \beta^\top U_i^\top \Sigma_i^{(w)-1} U_i \beta | \mathcal{D}_n] + \text{constants} \\ &= - (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n])^\top \Sigma_i^{(w)-1} U_i \beta + \frac{1}{2} \beta^\top U_i^\top \Sigma_i^{(w)-1} U_i \beta + \text{constants} \end{aligned}$$

(where we treat as constants all quantities independent of  $\beta$ ). Then, the gradient of  $A_i^3$  writes

$$\frac{\partial A_i^3(\beta, \Sigma_i^{(w)})}{\partial \beta} = - (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n])^\top \Sigma_i^{(w)-1} U_i + \beta^\top U_i^\top \Sigma_i^{(w)-1} U_i.$$

Given the form of  $U_i$  and  $V_i$  along with the fact that  $\Sigma_i^{(w)}$  is a diagonal matrix, we can rewrite the gradient of  $A_i^3$  as

$$\frac{\partial A_i^3(\beta, \Sigma_i^{(w)})}{\partial \beta} = - (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n])^\top U_i \tilde{\Sigma}^{(w)-1} + \beta^\top U_i^\top U_i \tilde{\Sigma}^{(w)-1},$$

where  $\tilde{\Sigma}^{(w)}$  is the diagonal matrix whose diagonal is  $(\phi_1^{(w)} \mathbf{1}_{q_1}^\top, \dots, \phi_L^{(w)} \mathbf{1}_{q_L}^\top)^\top \in \mathbb{R}^q$ . The closed-form update of  $\beta_k^{(w)}$  is then obtained by canceling the gradient, that is

$$\begin{aligned} &-n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \frac{\partial A_i^3(\beta, \Sigma_i^{(w)})}{\partial \beta} = 0 \\ \Leftrightarrow &\sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n])^\top U_i \tilde{\Sigma}^{(w)-1} - \beta^\top U_i^\top U_i \tilde{\Sigma}^{(w)-1} \right) = 0 \\ \Leftrightarrow &\sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n])^\top U_i - \beta^\top U_i^\top U_i \right) = 0 \\ \Leftrightarrow &\beta = \left( \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} U_i^\top U_i \right)^{-1} \left( \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} U_i^\top (Y_i - V_i \mathbb{E}_{\theta^{(w)}} [b_i | \mathcal{D}_n]) \right). \end{aligned}$$

LEMMA 5: *At step  $w+1$  of the EM algorithm, for any  $j \in \{1, \dots, J\}$ , given the the update of  $\lambda_0(\tau_j)$  is*

$$\lambda_0^{(w+1)}(\tau_j) = \frac{\sum_{i=1}^n \Delta_i \mathbb{1}_{\{T_i = \tau_j\}}}{\sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \exp(\psi_i(\tau_j)^\top \gamma_k^{(w+1)}) \mathbb{1}_{\{T_i \geq \tau_j\}}}, \quad (13)$$

where  $\tilde{\pi}_{ik}^{\theta^{(w)}}$  is given by (8).

*Proof.* The update for  $\lambda_0^{(w)}(\tau_j)$  requires to solve the minimization problem

$$\lambda_0^{(w+1)}(\tau_j) \in \operatorname{argmin}_{\lambda_0 \in \mathbb{R}} -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} A_i^4(\gamma_k^{(w+1)}, \lambda_0). \quad (14)$$

We have

$$\begin{aligned} &-n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} A_i^4(\gamma_k^{(w+1)}, \lambda_0) \\ &= -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \Delta_i (\log \lambda_0(T_i) + \psi_i(T_i)^\top \gamma_k^{(w+1)}) - \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k^{(w+1)}) \mathbb{1}_{\{\tau_j \leq T_i\}} \right) \\ &= -n^{-1} \sum_{i=1}^n \Delta_i \log \lambda_0(T_i) + n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k^{(w+1)}) \mathbb{1}_{\{\tau_j \leq T_i\}} + \text{constants} \end{aligned}$$

(where we keep only the terms with  $\tau_j$  for any  $j \in \{1, \dots, J\}$ ). By taking the gradient of the previous expression over  $\lambda_0(\tau_j)$  and setting it to zero, that is

$$-\frac{n^{-1}}{\lambda_0} \sum_{i=1}^n \Delta_i \mathbf{1}_{\{T_i = \tau_j\}} + n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \exp(\psi_i(\tau_j)^\top \gamma_k^{(w+1)}) \mathbf{1}_{\{T_i \geq \tau_j\}} = 0. \quad (15)$$

We then obtain the update for  $\lambda_0^{(w)}(\tau_j)$  as the desired result.

Note that the closed-form update of  $\lambda_0^{(w)}(\tau_j)$  is a Breslow-like estimator (Breslow, 1972) adapted to our model. Finally, recall that  $\Sigma_i$  is a diagonal matrix with a diagonal of the form  $(\phi_1 \mathbf{1}_{n_i^1}^\top \cdots \phi_L \mathbf{1}_{n_i^L}^\top)$ . Estimating  $\Sigma_i$  thus amounts to estimating  $\phi_1, \dots, \phi_L$ , whose updates are given in the following lemma.

LEMMA 6: *At step  $w + 1$  of the EM algorithm, the update of  $\phi_\ell$  is*

$$\begin{aligned} \phi_\ell^{(w+1)} = \frac{1}{\sum_{i=1}^n n_i^\ell} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} & \left( (Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)})^\top (Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)} - 2V_i^\ell \mathbb{E}_{\theta^{(w)}}[b_i^\ell | \mathcal{D}_n]) \right. \\ & \left. + \text{Tr}(V_i^{\ell\top} V_i^\ell \mathbb{E}_{\theta^{(w)}}[b_i^\ell b_i^{\ell\top} | \mathcal{D}_n]) \right), \end{aligned} \quad (16)$$

where  $\mathbb{E}_{\theta^{(w)}}[b_i^\ell | \mathcal{D}_n]$  and  $\mathbb{E}_{\theta^{(w)}}[b_i^\ell b_i^{\ell\top} | \mathcal{D}_n]$  are obtained from (6) and (7).

*Proof.* The update for  $\phi^{(w)}$  requires to solve the following minimization problem

$$\phi^{(w+1)} \in \underset{\phi \in \mathbb{R}^L}{\text{argmin}} -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} A_i^3(\beta_k^{(w+1)}, \text{Diag}_i(\phi)), \quad (17)$$

where we denote by  $\text{Diag}_i(\phi)$  the diagonal matrix  $\Sigma_i$  to make clear its dependence on  $\phi$ . We let  $M_{ik}^{(w+1)} = U_i \beta_k^{(w+1)} + V_i b_i$ ,  $M_{ik}^{\ell(w+1)} = U_i^\ell \beta_k^{\ell(w+1)} + V_i^\ell b_i^\ell$ . Then, taking advantage of the structure of  $\text{Diag}_i(\phi)$ , we have

$$\begin{aligned} A_i^3(\beta_k^{(w+1)}, \text{Diag}_i(\phi)) &= \frac{1}{2} \log \det(\text{Diag}_i(\phi)) + \mathbb{E}_{\theta^{(w)}}[(Y_i - M_{ik}^{(w+1)})^\top \text{Diag}_i(\phi)^{-1} (Y_i - M_{ik}^{(w+1)}) | \mathcal{D}_n] \\ &= \frac{1}{2} \sum_{\ell=1}^L n_i^\ell \log \phi_\ell + \sum_{\ell=1}^L \frac{1}{\phi_\ell} \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - M_{ik}^{\ell(w+1)})^\top (Y_i^\ell - M_{ik}^{\ell(w+1)}) | \mathcal{D}_n]. \end{aligned}$$

Then, the gradient of  $A_i^3$  along  $\phi_\ell$  is simply

$$\frac{\partial A_i^3(\beta_k^{(w+1)}, \text{Diag}_i(\phi))}{\partial \phi_\ell} = \frac{n_i^\ell}{\phi_\ell} - \frac{1}{\phi_\ell^2} \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - M_{ik}^{\ell(w+1)})^\top (Y_i^\ell - M_{ik}^{\ell(w+1)}) | \mathcal{D}_n],$$

The closed-form update of  $\phi_\ell^{(w)}$  is then obtained by setting

$$\begin{aligned} & -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \frac{\partial A_i^3(\beta_k^{(w+1)}, \text{Diag}_i(\phi))}{\partial \phi_\ell} = 0 \\ \Leftrightarrow & -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \frac{n_i^\ell}{\phi_\ell} - \frac{1}{\phi_\ell^2} \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - M_{ik}^{\ell(w+1)})^\top (Y_i^\ell - M_{ik}^{\ell(w+1)}) | \mathcal{D}_n] \right) = 0 \\ \Leftrightarrow & \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \left( n_i^\ell \phi_\ell^{(w+1)} - \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - M_{ik}^{\ell(w+1)})^\top (Y_i^\ell - M_{ik}^{\ell(w+1)}) | \mathcal{D}_n] \right) = 0. \end{aligned}$$

The result follows from the fact that

$$\begin{aligned} & \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - M_{ik}^{\ell(w+1)})^\top (Y_i^\ell - M_{ik}^{\ell(w+1)}) | \mathcal{D}_n] \\ &= \mathbb{E}_{\theta^{(w)}}[(Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)})^\top (Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)} - 2V_i^\ell b_i^\ell) + b_i^{\ell\top} V_i^{\ell\top} V_i^\ell b_i | \mathcal{D}_n] \\ &= (Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)})^\top (Y_i^\ell - U_i^\ell \beta_k^{\ell(w+1)} - 2V_i^\ell \mathbb{E}_{\theta^{(w)}}[b_i^\ell | \mathcal{D}_n]) + \text{Tr}(V_i^{\ell\top} V_i^\ell \mathbb{E}_{\theta^{(w)}}[b_i^\ell b_i^{\ell\top} | \mathcal{D}_n]). \end{aligned}$$

A.3 *M-step: Update  $\xi$* 

In  $\mathcal{Q}_n(\theta, \theta^{(w)})$ , the parameter  $\xi$  appears only in the term

$$\begin{aligned} & -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} A_{i,k}^2(\xi) \\ &= -n^{-1} \sum_{i=1}^n \sum_{k=1}^K \tilde{\pi}_{ik}^{\theta^{(w)}} \log \left( \frac{e^{X_i^\top \xi_k}}{\sum_{j=1}^K e^{X_i^\top \xi_j}} \right) \\ &= -n^{-1} \sum_{i=1}^n \left( \tilde{\pi}_{ik}^{\theta^{(w)}} \log \left( 1 + \sum_{\substack{j \neq k \\ j=1}}^K e^{X_i^\top (\xi_j - \xi_k)} \right) + \sum_{\substack{m \neq k \\ m=1}}^K \tilde{\pi}_{im}^{\theta^{(w)}} \log \left( 1 + e^{X_i^\top (\xi_k - \xi_m)} + \sum_{\substack{j \neq k, j \neq m \\ j=1}}^K e^{X_i^\top (\xi_j - \xi_m)} \right) \right). \end{aligned}$$

For  $k \in \{1, \dots, K\}$ , the update for  $\xi_k^{(w)}$  requires to solve the minimization problem

$$\xi_k^{(w+1)} \in \underset{\xi \in \mathbb{R}^p}{\operatorname{argmin}} \mathcal{F}_{1,k}(\xi) + \zeta_{1,k} \Omega_1(\xi), \quad (18)$$

where  $\mathcal{F}_{1,k}$  is defined by

$$\begin{aligned} \mathcal{F}_{1,k}(\xi) &= n^{-1} \sum_{i=1}^n \left( \tilde{\pi}_{ik}^{\theta^{(w)}} \log \left( 1 + \sum_{\substack{j \neq k \\ j=1}}^K e^{X_i^\top (\xi_j - \xi)} \right) \right. \\ &\quad \left. + \sum_{\substack{m \neq k \\ m=1}}^K \tilde{\pi}_{im}^{\theta^{(w)}} \log \left( 1 + e^{X_i^\top (\xi - \xi_m)} + \sum_{\substack{j \neq k, j \neq m \\ j=1}}^K e^{X_i^\top (\xi_j - \xi_m)} \right) \right) \end{aligned}$$

and  $\Omega_1$  is the elastic net regularization. We choose to solve (18) using the L-BFGS-B algorithm (Zhu et al., 1997) which belongs to the class of quasi-Newton optimization routines and solves the given minimization problem by computing approximations of the inverse Hessian matrix of the objective function. It can deal with differentiable convex objectives with box constraints.

In order to use it with  $l_1$  part of the elastic net regularization, which is not differentiable, we use the trick borrowed from Andrew and Gao (2007): for  $a \in \mathbb{R}$ , write  $|a| = a^+ + a^-$ , where  $a^+$  and  $a^-$  are respectively the positive and negative part of  $a$ , and add the constraints  $a^+ \geq 0$  and  $a^- \geq 0$ . Namely, we rewrite the minimization problem (18) as the following differentiable problem with box constraints

$$\begin{aligned} & \text{minimize} \quad \mathcal{F}_{1,k}(\xi^+ - \xi^-) + \zeta_{1,k} \left( (1 - \eta) \sum_{j=1}^p (\xi_j^+ + \xi_j^-) + \frac{\eta}{2} \|\xi^+ - \xi^-\|_2^2 \right) \\ & \text{subject to} \quad \xi_j^+ \geq 0 \text{ and } \xi_j^- \geq 0 \text{ for } j \in \{1, \dots, p\} \end{aligned} \quad (19)$$

where  $\xi^\pm = (\xi_1^\pm, \dots, \xi_p^\pm)^\top$ . The L-BFGS-B solver requires the exact value of the gradient, which is easily given by

$$\frac{\partial \mathcal{F}_{1,k}(\xi)}{\partial \xi} = -n^{-1} \sum_{i=1}^n \left( \tilde{\pi}_{ik}^{\theta^{(w)}} - \frac{e^{X_i^\top \xi}}{e^{X_i^\top \xi} + \sum_{\substack{j \neq k \\ j=1}}^K e^{X_i^\top \xi_j}} \right) X_i^\top. \quad (20)$$

In practice, we use the Python solver `fmin_l_bfgs_b` from `scipy.optimize` (Virtanen et al., 2020).

A.4 *M-step: Update  $\gamma$* 

In  $\mathcal{Q}_n(\theta, \theta^{(w)})$ , for  $k \in \{1, \dots, K\}$ ,  $\gamma_k$  appears only in the term

$$\begin{aligned} & -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} A_i^4(\gamma_k, \lambda_0) \\ &= -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \Delta_i (\log \lambda_0(T_i) + \psi_i(T_i)^\top \gamma_k) - \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k) \mathbf{1}_{\{\tau_j \leq T_i\}} \right) \\ &= -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \Delta_i \psi_i(T_i)^\top \gamma_k - \sum_{j=1}^J \lambda_0(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma_k) \mathbf{1}_{\{\tau_j \leq T_i\}} \right) + \text{constants}. \end{aligned}$$

Then the update for  $\gamma_k^{(w)}$  requires to solve the following minimization problem

$$\gamma_k^{(w+1)} \in \underset{\gamma \in \mathbb{R}^{LM}}{\operatorname{argmin}} \mathcal{F}_{2,k}(\gamma) + \zeta_{2,k} \Omega_2(\gamma), \quad (21)$$

where  $\mathcal{F}_{2,k}$  is defined by

$$\mathcal{F}_{2,k}(\gamma) = -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \Delta_i \psi_i(T_i)^\top \gamma - \sum_{j=1}^J \lambda_0^{(w)}(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma) \mathbf{1}_{\{\tau_j \leq T_i\}} \right)$$

and  $\Omega_2$  is the sparse group lasso regularization. We choose to solve problem (21) using the iterative soft-thresholding algorithm (ISTA), which is a proximal gradient descent algorithm (Beck and Teboulle, 2009). In our context, this method requires the gradient of  $\mathcal{F}_{2,k}$  as well as the proximal operator (Moreau, 1962) of the sparse group lasso. We refer to the proof of Yuan et al. (2011, Theorem 1) to show that the proximal operator of the sparse group lasso can be expressed as the composition of the group lasso and the lasso proximal operators, which are both well known analytically (Bach et al., 2012) and tractable. The gradient of  $\mathcal{F}_{2,k}$  is here given by

$$\frac{\partial \mathcal{F}_{2,k}(\gamma)}{\partial \gamma} = -n^{-1} \sum_{i=1}^n \tilde{\pi}_{ik}^{\theta^{(w)}} \left( \Delta_i - \sum_{j=1}^J \lambda_0^{(w)}(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma) \mathbf{1}_{\{T_i \geq \tau_j\}} \right) \psi_i(T_i)^\top. \quad (22)$$

We use the `Python` library `copt` for the implementation of proximal gradient descent (Fabian Pedregosa, 2020) and we propose in our `FLASH` package a first `Python` implementation for the proximal operator of the sparse group lasso.

#### A.5 Convex optimization problems with respect to $\xi$ and $\gamma$

LEMMA 7: *The optimization problems defined in (18) and (21) are convex.*

*Proof.* Given first order derivative of  $\mathcal{F}_{1,k}(\xi)$  in (20), we show that the second order derivative of  $\mathcal{F}_{1,k}(\xi)$  is positive definite

$$\frac{\partial^2 \mathcal{F}_{1,k}(\xi)}{\partial \xi \partial \xi^\top} = n^{-1} \sum_{i=1}^n \left( \frac{e^{X_i^\top \xi}}{e^{X_i^\top \xi} + \sum_{\substack{j \neq k \\ j=1}}^K e^{X_i^\top \xi_j}} \right) \left( 1 - \frac{e^{X_i^\top \xi}}{e^{X_i^\top \xi} + \sum_{\substack{j \neq k \\ j=1}}^K e^{X_i^\top \xi_j}} \right) X_i X_i^\top \in S_{++}^p.$$

Given first order derivative of  $\mathcal{F}_{2,k}(\gamma)$  in (22), we show that the second order derivative of  $\mathcal{F}_{2,k}(\gamma)$  is positive definite

$$\frac{\partial^2 \mathcal{F}_{2,k}(\gamma)}{\partial \gamma \partial \gamma^\top} = n^{-1} \sum_{i=1}^n \left( \tilde{\pi}_{ik}^{\theta^{(w)}} \sum_{j=1}^J \lambda_0^{(w)}(\tau_j) \exp(\psi_i(\tau_j)^\top \gamma) \mathbf{1}_{\{T_i \geq \tau_j\}} \right) \psi_i(T_i) \psi_i(T_i)^\top \in S_{++}^{LM}.$$

As we already defined the elastic net and sparse group lasso regularization

$$\Omega_1(\xi) = (1 - \eta) \|\xi\|_1 + \frac{\eta}{2} \|\xi\|_2^2 \quad \text{and} \quad \Omega_2(\gamma) = (1 - \tilde{\eta}) \|\gamma\|_1 + \tilde{\eta} \sum_{\ell=1}^L \|\gamma^\ell\|_2,$$

with  $(\eta, \tilde{\eta}) \in [0, 1]^2$ . Note that every norm is convex and a non-negative weighted sum of convex functions is convex (Boyd and Vandenberghe, 2004) then  $\Omega_1(\xi)$  and  $\Omega_2(\gamma)$  are convex functions. Therefore,  $\mathcal{F}_{1,k}(\xi) + \Omega_1(\xi)$  and  $\mathcal{F}_{2,k}(\gamma) + \Omega_2(\gamma)$  are strictly convex.

#### A.6 The extended EM algorithm

Algorithm 1 below describes the main steps of our proposed EM algorithm.

#### A.7 Monotone convergence

By denoting  $\zeta_p^{(w)} = (\theta_1^{(w+1)}, \dots, \theta_p^{(w+1)}, \theta_{p+1}^{(w)}, \dots, \theta_p^{(w)})$  and from the definition of our proposed EM algorithm, at step  $w + 1$ , we have

$$\mathcal{Q}_n^{\text{pen}}(\theta^{(w)}, \theta^{(w)}) \geq \mathcal{Q}_n^{\text{pen}}(\zeta_1^{(w)}, \theta^{(w)}) \geq \dots \geq \mathcal{Q}_n^{\text{pen}}(\zeta_{p-1}^{(w)}, \theta^{(w)}) \geq \mathcal{Q}_n^{\text{pen}}(\theta^{(w+1)}, \theta^{(w)}). \quad (23)$$

Then, we are in a generalized EM (GEM) setting (Dempster et al., 1977), where

$$\mathcal{Q}_n^{\text{pen}}(\theta^{(w+1)}, \theta^{(w)}) \leq \mathcal{Q}_n^{\text{pen}}(\theta^{(w)}, \theta^{(w)}). \quad (24)$$

For such algorithms, we refer to monotonicity of the likelihood property from the book of McLachlan and Krishnan (2007, Section 3.3) to show that the objective function (1) decreases at each iteration, namely

$$\mathcal{L}_n^{\text{pen}}(\theta^{(w+1)}) \leq \mathcal{L}_n^{\text{pen}}(\theta^{(w)}).$$

---

**Algorithm 1** The extended EM algorithm for FLASH inference

---

**Data:** Training data  $\mathcal{D}_n$ ; tuning hyper-parameters  $(\zeta_{1,k}, \zeta_{2,k})_{k \in \{1, \dots, K\}}$   
**Input:** maximum iteration  $W$ , tolerance  $\varepsilon$   
**Output:** Last parameters  $\hat{\theta} \in \mathbb{R}^P$

- 1: Initialize parameters  $\theta^{(0)} \in \mathbb{R}^P$
- 2: **for**  $w = 1, \dots, W$  **do**
  - E-step:**
  - 3: Compute  $(\mathbb{E}_{\theta^{(w)}}[b_i | T_i, \Delta_i, Y_i])_{i \in \{1, \dots, n\}}$ ,  $(\mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i])_{i \in \{1, \dots, n\}}$ ,  $(\tilde{\pi}_{ik}^{\theta^{(w)}})_{\substack{i \in \{1, \dots, n\} \\ k \in \{1, \dots, K\}}}$
  - 4: Compute  $(\mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | T_i, \Delta_i, Y_i])_{i \in \{1, \dots, n\}}$
  - 5: Compute  $(\tilde{\pi}_{ik}^{\theta^{(w)}})_{\substack{i \in \{1, \dots, n\} \\ k \in \{1, \dots, K\}}}$
  - M-step:**
  - 6: Update  $D^{(w+1)}$
  - 7: Update  $(\xi_k^{(w+1)})_{k \in \{1, \dots, K\}}$  with L-BFGS-B
  - 8: Update  $(\beta_k^{(w+1)})_{k \in \{1, \dots, K\}}$
  - 9: Update  $(\gamma_k^{(w+1)})_{k \in \{1, \dots, K\}}$  with proximal gradient descent
  - 10: Update  $\lambda_0^{(w+1)}$  and  $\phi^{(w+1)}$
  - 11: **if**  $(\mathcal{L}_n^{\text{pen}}(\theta^{(w+1)}) - \mathcal{L}_n^{\text{pen}}(\theta^{(w)})) / \mathcal{L}_n^{\text{pen}}(\theta^{(w)}) < \varepsilon$  **then**  
    - break**
  - 12: **end if**
- 13: **end for**
- 14: **Return**  $\hat{\theta} = \theta^{(w+1)}$

---

## B. MATHEMATICAL DETAILS OF JLCMS AND SREMS

Given our notation used in the main body of the paper, we define here the sub-models of the two main approaches of joint models: JLCMs and SREMs.

*JLCMs.* It assumes that the population is heterogeneous and that there are homogeneous latent classes that share the same marker trajectories and the same prognosis. The latent class membership probability is assumed to take the form of multinomial logistic regression

$$\mathbb{P}[g_i = k | X_i] = \frac{e^{X_i^\top \xi_k}}{\sum_{j=1}^K e^{X_i^\top \xi_j}}.$$

The dependence between the time-to-event and the longitudinal marker is fully captured by a latent class structure. There are no shared associations between the longitudinal and survival models. Given the latent class membership, each submodel is assumed to be independent. If we choose Gaussian linear model for longitudinal markers and Cox relative risk model for the time-to-event, we have

$$y_i^\ell(t_{ij}^\ell) | b_i^\ell, g_i = k \sim \mathcal{N}(m_{ik}^\ell(t_{ij}^\ell), \phi_\ell) \quad \text{and} \quad \lambda(t | g_i = k) = \lambda_0(t) \exp\left(X_i^\top \gamma_k\right),$$

where  $\gamma_k$  is the  $p$ -vector of unknown parameters. We consider the implementation of JLCMs in R package LCMM (function `mpjlcmm`). In this context, the predictive marker for subject  $i$  at time  $s_i$  is

$$\hat{\mathcal{R}}_{ik}(s_i) = \frac{\mathbb{P}_{\hat{\theta}}(g_i = k) f_{\hat{\theta}}(T_i = s_i, \Delta_i = 0, \mathcal{Y}_i(s_i^-) | g_i = k)}{\sum_{j=1}^K \mathbb{P}_{\hat{\theta}}(g_i = j) f_{\hat{\theta}}(T_i = s_i, \Delta_i = 0, \mathcal{Y}_i(s_i^-) | g_i = j)},$$

where the density  $f_{\hat{\theta}}$  are the one corresponding to a JLCM model.

*SREMs.* It assumes a homogeneous population of subjects and the dependency between the time-to-event and the longitudinal marker is influenced by some random effects learned in a linear mixed model. If we choose Gaussian linear model for longitudinal markers, we have

$$y_i^\ell(t_{ij}^\ell) | b_i^\ell \sim \mathcal{N}(m_i^\ell(t_{ij}^\ell), \phi_\ell),$$

where  $m_i^\ell(t_{ij}^\ell) = u^\ell(t_{ij}^\ell)^\top \beta^\ell + v^\ell(t_{ij}^\ell)^\top b_i^\ell$  and  $\beta^\ell$  is a  $r_\ell$ -vector of unknown fixed effect parameters. The random effects are included as covariates in the survival model through the shared association functions  $\phi$ . If we choose Cox relative risk model for the time-to-event, we have

$$\lambda(t | g_i = k) = \lambda_0(t) \exp\left(X_i^\top \gamma_0 + \sum_{\ell=1}^L \phi(b_i^\ell, t_i)^\top \gamma^\ell\right),$$



where  $\gamma_0$  and  $\gamma^\ell$  are unknown parameters.

We consider the implementation of SREMs in R package `JMbayes`. In this context, the predictive marker for subject  $i$  at time  $s_i$  is

$$\widehat{\mathcal{R}}_i(s_i) = \exp\left(X_i^\top \gamma_0 + \sum_{\ell=1}^L \phi(b_i^\ell, s_i)^\top \gamma^\ell\right),$$

where the shared associations takes the form  $\phi(b_i^\ell, s_i) = u^\ell(s_i)^\top \beta^\ell + v^\ell(s_i)^\top b_i^\ell$ .

## C. EXPERIMENTAL DETAILS AND ADDITIONAL EXPERIMENTS

### C.1 Initialization

**C.1.1 Initialization.** In order to help convergence,  $\theta^{(0)}$  should be well chosen. We then give some details about the starting point  $\theta^{(0)}$  of this algorithm. For all  $k = 1, \dots, K$ , we first choose  $\xi_k^{(0)} = \mathbf{0}_d$  and  $\gamma_k^{(0)} = 0.01 * \mathbf{1}_{LA}$ . Then, we initialize  $\lambda_0^{(0)}$  like if there is no latent classes ( $\gamma_1^{(0)} = \dots = \gamma_K^{(0)}$ ) with a standard Cox proportional hazards regression with time-independent features. Finally, the longitudinal submodel parameters  $\beta_k^{(0)}$ ,  $D^{(0)}$  and  $\phi^{(0)}$  are initialized – again like if there is no latent classes ( $\beta_1^{(0)} = \dots = \beta_K^{(0)}$ ) – using a multivariate linear mixed model with an explicit EM algorithm, being itself initialized with univariate fits.

**C.1.2 Multivariate linear mixed model.** Let us derive here the explicit EM algorithm for the multivariate Gaussian linear mixed model used to initialize the longitudinal parameters  $\beta_k^{(0)}$ ,  $D^{(0)}$  and  $\phi^{(0)}$  in the proposed EM algorithm in Section C.1, acting as if there is no latent classes ( $\beta_1^{(0)} = \dots = \beta_K^{(0)}$ ). For the sake of simplicity, let us denote here

$$\theta = (\beta^\top, D, \phi^\top)^\top \in \mathbb{R}^P$$

the parameter vector to infer. The conditional distribution of  $Y_i|b_i$  then writes

$$f(Y_i|b_i; \theta) = -(2\pi)^{-\frac{n_i}{2}} \det(\Sigma_i)^{-\frac{1}{2}} \exp^{-\frac{1}{2}(Y_i - M_i)^\top \Sigma_i^{-1}(Y_i - M_i)},$$

where  $M_i = U_i\beta + V_i b_i$  in this context. The negative complete log-likelihood then writes

$$\begin{aligned} \mathcal{L}_n^{\text{comp}}(\theta) &= \mathcal{L}_n^{\text{comp}}(\theta; \mathcal{D}_n, \mathbf{b}) \\ &= \sum_{i=1}^n -\frac{1}{2} (n_i \log 2\pi + \log \det(\Sigma_i) + (Y_i - M_i)^\top \Sigma_i^{-1} (Y_i - M_i)) \\ &\quad - \frac{1}{2} (r \log 2\pi + \log \det(D) + b_i^\top D^{-1} b_i). \end{aligned}$$

*E-step.* Supposing that we are at step  $w + 1$  of the algorithm, with current iterate denoted  $\theta^{(w)}$ , we need to compute the expectation of the negative complete log-likelihood conditional on the observed data and the current estimate of the parameters, which is given by

$$\mathcal{Q}_n(\theta, \theta^{(w)}) = \mathbb{E}_{\theta^{(w)}}[\mathcal{L}_n^{\text{comp}}(\theta) | \mathcal{D}_n].$$

Here, the calculation of this quantity is reduced to the calculation of  $\mathbb{E}_{\theta^{(w)}}[b_i|Y_i]$  and  $\mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | Y_i]$  for  $i = 1, \dots, n$ . The marginal distributions of  $Y_i$  and  $b_i$  being both Gaussian, one has from Bayes Theorem

$$f(b_i|Y_i; \theta^{(w)}) \propto \exp\left(-\frac{1}{2}(b_i - \mu_i^{(w)})^\top \Omega_i^{(w)-1}(b_i - \mu_i^{(w)})\right)$$

where

$$\Omega_i^{(w)} = (V_i^\top \Sigma_i^{(w)-1} V_i + D^{(w)-1})^{-1} \quad \text{and} \quad \mu_i^{(w)} = \Omega_i^{(w)} V_i^\top \Sigma_i^{(w)-1} (Y_i - U_i \beta^{(w)}).$$

Then, one has

$$\begin{cases} \mathbb{E}_{\theta^{(w)}}[b_i|Y_i] = \mu_i^{(w)}, \\ \mathbb{E}_{\theta^{(w)}}[b_i b_i^\top | Y_i] = \Omega_i^{(w)} + \mu_i^{(w)} \mu_i^{(w)\top}. \end{cases}$$

*M-step.* Here, we need to compute

$$\theta^{(w+1)} \in \operatorname{argmin}_{\theta \in \mathbb{R}^P} \mathcal{Q}_n(\theta, \theta^{(w)}).$$

The parameters updates are then naturally given in closed form by zeroing the gradient. One obtains

$$\beta^{(w+1)} = \left( \sum_{i=1}^n U_i^\top U_i \right)^{-1} \left( \sum_{i=1}^n U_i^\top Y_i - U_i V_i \mathbb{E}_{\theta^{(w)}} [b_i | Y_i] \right),$$

$$\phi_\ell^{(w+1)} = \left( \sum_{i=1}^n n_i^\ell \right)^{-1} \left( \sum_{i=1}^n (Y_i^\ell - U_i^\ell \beta_\ell^{(w+1)})^\top (Y_i^\ell - U_i^\ell \beta_\ell^{(w+1)} - 2V_i^\ell \mathbb{E}_{\theta^{(w)}} [b_i^\ell | Y_i^\ell]) \right. \\ \left. + \text{Tr} (V_i^{\ell \top} V_i^\ell \mathbb{E}_{\theta^{(w)}} [b_i^\ell b_i^{\ell \top} | Y_i^\ell]) \right)$$

and

$$D^{(w+1)} = n^{-1} \sum_{i=1}^n \mathbb{E}_{\theta^{(w)}} [b_i b_i^\top | Y_i].$$

**C.1.3 Implementation of multivariate linear mixed model.** We implement an EM algorithm for fitting a multivariate linear mixed model used to initialize parameters of longitudinal submodel in Algorithm 1. Let us introduce the list  $\Omega^{(w)} = [\Omega_1^{(w)}, \dots, \Omega_n^{(w)}]$ , the matrices  $\mu = [\mu_1, \dots, \mu_n] \in \mathbb{R}^{r \times n}$ ,  $U^\ell = [U_1^{\ell \top} \dots U_n^{\ell \top}]^\top \in \mathbb{R}^{n_\ell \times q_\ell}$ ,  $U = [U_1^\top \dots U_n^\top]^\top \in \mathbb{R}^{\mathcal{N} \times q}$ ,

$$V^\ell = \begin{bmatrix} V_1^\ell & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & V_n^\ell \end{bmatrix}, \quad V = \begin{bmatrix} V_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & V_n \end{bmatrix}, \quad \Omega^{\ell(w)} = \begin{bmatrix} \Omega_1^{\ell(w)} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Omega_n^{\ell(w)} \end{bmatrix}$$

that belong respectively in  $\mathbb{R}^{n_\ell \times n r_\ell}$ ,  $\mathbb{R}^{\mathcal{N} \times n r}$  and  $\mathbb{R}^{n r_\ell \times n r_\ell}$ , and the vectors  $\tilde{\mu}^{(w)} = (\mu_1^{(w)\top} \dots \mu_n^{(w)\top})^\top \in \mathbb{R}^{n r}$ ,  $(\tilde{\mu}^\ell)^{(w)} = ((\mu_1^\ell)^{(w)\top} \dots (\mu_n^\ell)^{(w)\top})^\top \in \mathbb{R}^{n r_\ell}$ ,  $y^\ell = (y_1^{\ell \top} \dots y_n^{\ell \top})^\top \in \mathbb{R}^{n_\ell}$  with  $n_\ell = \sum_{i=1}^n n_i^\ell$  and  $y = (y_1^\top \dots y_n^\top)^\top \in \mathbb{R}^{\mathcal{N}}$  with  $\mathcal{N} = \sum_{i=1}^n n_i$ . The  $\beta$  update then rewrites

$$\beta^{(w+1)} = (U^\top U)^{-1} U^\top (y - V \tilde{\mu}^{(w)}).$$

For the  $D$  update, one has

$$D^{(w+1)} = n^{-1} (\text{sum}(\Omega^{(w)}) + \mu^{(w)} \mu^{(w)\top}).$$

And finally for the  $\phi$  update, one has

$$\phi_\ell^{(w+1)} = n_\ell^{-1} [(y^\ell - U^\ell \beta_\ell^{(w+1)})^\top (y^\ell - U^\ell \beta_\ell^{(w+1)} - 2V^\ell (\tilde{\mu}^\ell)^{(w)}) \\ + \text{Tr}\{V^{\ell \top} V^\ell (\Omega^{\ell(w)} + (\tilde{\mu}^\ell)^{(w)} (\tilde{\mu}^\ell)^{(w)\top})\}].$$

In our implementation, these parameters are initialized with univariate fits using the function `mixedlm` (linear mixed effects model) in the Python package `statsmodels`.

### C.2 Details of the simulation setting

We assume that each of the  $n$  subjects belongs to one of two different profiles: high-risk and low-risk. Let us denote by  $\mathcal{H} \subset \{1, \dots, n\}$  the set of high-risk subjects. For generating the time-independent features matrix of subject  $i$ , we take

$$X_i \in \mathbb{R}^p \sim \begin{cases} \mathcal{N}(\mu, \Sigma_1(\rho_1)) & \text{if } i \notin \mathcal{H}, \\ \mathcal{N}(-\mu, \Sigma_1(\rho_1)) & \text{if } i \in \mathcal{H}, \end{cases}$$

where the mean  $\mu$  corresponds to the gap between time-independent features of high-risk subjects and low-risk subjects, and  $\Sigma_1(\rho_1)$  a  $p \times p$  Toeplitz covariance matrix (Mukherjee and Maiti, 1988) with correlation  $\rho_1 \in (0, 1)$ , that is,  $\Sigma_1(\rho_1)_{jj'} = \rho_1^{|j-j'|}$ . In order to simulate the class  $g_i$  for each subject  $i$ , we choose a sparse coefficient vector where we decide to keep only  $\bar{p}$  active features, that is

$$\xi = (\nu, \dots, \nu, 0, \dots, 0) \in \mathbb{R}^p, \quad (25)$$

with  $\nu \in \mathbb{R}$  being the value of the active coefficients. Then, we generate  $g_i \sim \mathcal{B}(\pi_\xi(X_i))$ , where  $\mathcal{B}(\alpha)$  denotes the Bernoulli distribution with parameter  $\alpha \in [0, 1]$  and

$$\pi_\xi(X_i) = \frac{e^{X_i^\top \xi}}{1 + e^{X_i^\top \xi}}$$

Now, concerning the simulation of longitudinal markers, the idea is to sample from multivariate normal distributions.

Moreover, we want to induce sparsity in the longitudinal data to reduce correlation between longitudinal features in each class  $k$ . We denote by  $\mathcal{S}_k$  the set of active longitudinal features in class  $k$ , which is randomly selected from the set  $\{1, \dots, L\}$ . Then, we simulate longitudinal features of the form

$$Y_i^\ell(t) = \sum_{k=1}^K \mathbb{1}_{\{g_i=k\}} \left( ((1, t)^\top \beta_k^\ell + (1, t)^\top b_k^\ell) \mathbb{1}_{\{\ell \in \mathcal{S}_k\}} + \varepsilon_i^\ell(t) \right)$$

where  $t \geq 0$ , the error term  $\varepsilon_i^\ell(t) \sim \mathcal{N}(0, \sigma_\ell^2)$ , the global variance-covariance matrix for the random effects components is such that  $D = \Sigma_2(\rho_2)$ , a  $r \times r$  Toeplitz covariance matrix with correlation  $\rho_2 \in (0, 1)$ , and the fixed effect parameters are generated according to

$$\beta_k^\ell \sim \mathcal{N}\left(\mu_k, \begin{bmatrix} \rho_3 & 0 \\ 0 & \rho_3 \end{bmatrix}\right)$$

for  $k \in \{1, 2\}$  and with correlation  $\rho_3 \in (0, 1)$ . The number of observations for each subject is randomly selected from 1 to 10, and the measurement times are simulated from a uniform distribution with minimum zero and maximum its survival time.

Now to generate survival times, we choose a risk model with

$$\lambda_i(t | g_i = k) = \lambda_0(t) \exp\left(\sum_{\ell=1}^L \Psi_{i,k}^\ell(t) \gamma_k^\ell\right), \quad (26)$$

We choose a Gompertz distribution (Gompertz, 1825) for the baseline, that is

$$\lambda_0(t) = \kappa_1 \kappa_2 \exp(\kappa_2 t) \quad (27)$$

with  $\kappa_1 > 0$  and  $\kappa_2 \in \mathbb{R}$  the scale and shape parameters respectively, which is a common distribution choice in survival analysis (Klein and Moeschberger, 2005) with a rich history in describing mortality curves. For the choice of the association features, we consider the two functionals in form of random effects linear predictor  $m_k^\ell(t)$  (Chi and Ibrahim, 2006) and random effects  $b^\ell$  (Hatfield et al., 2011), that is  $\Psi_{i,k}^\ell(t) = (\beta_{k,1}^\ell + \beta_{k,2}^\ell t + b_{i,1}^\ell + b_{i,2}^\ell t, b_i^{\ell\top})^\top$  and  $\gamma_k^\ell = \nu_k \mathbb{1}_{\{\ell \in \mathcal{S}_k\}}$  where  $\nu_k \in \mathbb{R}$  is the coefficients of active association features for each group  $k$ . Then one can write

$$\lambda_i(t | g_i = k) = \lambda_0(t) \exp(\iota_{i,k,1} + \iota_{i,k,2} t),$$

being a Cox model with a linear relationship between time-varying feature and log hazard that allows the following explicit survival times generation process. One can now generate survival times explicitly, via the inversion method, as

$$T_i^* | g_i = k \sim \frac{1}{\iota_{i,k,2} + \kappa_2} \log\left(1 - \frac{(\iota_{i,k,2} + \kappa_2) \log U_i}{\kappa_1 \kappa_2 \exp \iota_{i,k,1}}\right) \quad (28)$$

where  $U_i \sim \mathcal{U}([0, 1])$ , see Austin (2013). The distribution of the censoring variable  $C_i$  is the geometric distribution  $\mathcal{G}(\alpha_c)$ , where  $\alpha_c \in (0, 1)$  is empirically tuned to maintain a desired censoring rate  $r_c \in [0, 1]$ . The choice of all hyper-parameters is driven by the applications on simulated data presented in Section 5.1 in the main paper, and summarized in Table 3.

Table 3: Hyper-parameter choices for simulation with  $n = 500$ ,  $L = 5$  and  $p = 10$ .

$ \mathcal{H} $	$ \mathcal{S}_k $	$(\rho_1, \rho_2, \rho_3)$	$\mu$	$\mu_1$	$\mu_2$	$\sigma_\ell^2$	$(\kappa_1, \kappa_2)$	$(\nu, \nu_1, \nu_2)$	$r_c$	$\bar{p}$
200	2	(0.5, 0.01, 0.01)	1	$\begin{pmatrix} -0.6 \\ 0.1 \end{pmatrix}$	$\begin{pmatrix} 0.05 \\ 0.2 \end{pmatrix}$	0.25	(0.05, 0.1)	(0.2, 0.1, 0.4)	0.3	5

### C.3 Description of the datasets used in comparison study

*JoinerML simulation.* We use the classical R package `joinerML` (Hickey et al., 2018) to simulate multivariate longitudinal and time-to-event data from a joint model. The multivariate longitudinal features are generated for all possible measurement times using multivariate Gaussian linear mixed model. Failure times are simulated from proportional hazards time-to-event models. We sample two time-independent features and two longitudinal features for 250 individuals.

*PBCseq dataset.* This dataset which is available in the R package `JMbayes` (Rizopoulos, 2016), contains the follow-up of 312 patients with primary biliary cirrhosis, a rare autoimmune liver disease. Several longitudinal features are measured over time (for example serum bilirubin, serum cholesterol, albumin), along with information on gender, age, and drug used recorded once at the beginning of the study. Time-to-event is also recorded with a censoring rate of 55%.

*Aids dataset.* This dataset which is available in the R package `JMbayes` (Rizopoulos, 2016), compares the efficacy and safety of two drugs for 467 patients diagnosed with HIV who were either intolerant or resistant to zidovudine therapy. Information on gender, age, drug used, AIDS infection status, and level of intolerance to zidovudine is collected at the start of the study. The

longitudinal feature of interest here is the measurement of the number of CD4 cell (a type of white blood cell), a laboratory test used to understand the progression of HIV disease. Time-to-event is also recorded with a censoring rate of 40%.

#### C.4 Procedure to evaluate model performance

Let us describe now in Algorithm 2 the procedure we follow to evaluate model performance on simulated data and real data described in Section 5 in the main paper.

---

**Algorithm 2** Procedure followed to assess performances of a given model in our real-time prediction paradigm.

---

**Input:** Dataset  $\mathcal{D}_n$ ; a model under study

**Output:** Confidence intervals on C-index metric as well as on running time.

```

1: We run  $K^{\text{iter}} = 50$  independent experiments
2: for  $k = 1, \dots, K^{\text{iter}}$  do
3:   start_time = time()
4:    $(\mathcal{D}_{n_{\text{train}}}, \mathcal{D}_{n_{\text{test}}}) \leftarrow \text{split\_train\_test}(\mathcal{D}_n)$ 
5:   model.fit $(\mathcal{D}_{n_{\text{train}}})$ 
6:   for  $i = 1, \dots, n_{\text{test}}$  do
7:      $s_i \sim \max_{\ell \in \{1, \dots, L\}} (t_{i n_i}^\ell) \times (1 - \text{Beta}(2, 5))$ 
8:      $Y_i \leftarrow (Y_i^\ell(t_{ij}^\ell))_{j \in \{1, \dots, n_i^\ell - 1\}}$  with  $t_{ij}^\ell \leq s_i$ 
9:      $\mathcal{X}_i = (X_i, Y_i)$ 
10:     $\widehat{\mathcal{R}}_i^k(s_i) \leftarrow \text{model.predict}(\mathcal{X}_i)$ 
11:   end for
12:   score $_k \leftarrow \text{c\_index}((\widehat{\mathcal{R}}_i^k(t_i), T_i, \Delta_i)_{i=1, \dots, n_{\text{test}}})$ 
13:   end_time = time()
14:   running_time $_k = \text{end\_time} - \text{start\_time}$ 
15: end for
16: Return  $\hat{\theta} = \theta^{(w+1)}$ 

```

---

*Screening phase.* We use the multivariate Cox model and C-index metric for selecting the  $M$  most important features from a specific set of feature extraction functions  $\mathcal{F}$ . For each feature in  $\mathcal{F}$ , we extract the  $n \times L$  matrix from all  $L$  longitudinal markers of all  $n$  subjects. Then we fit this extracted matrix with all the survival times  $T$  and censoring indicators  $\Delta$  in the Cox model and use the C-index metric to evaluate the performance. Finally, we select  $M$  features have the highest C-index values.

#### C.5 Interpretation of the model on medical datasets

Tables 4a and 4b provide the estimated coefficients of FLASH on the PBCseq and Sepsis datasets. Coefficients are organized as follows: first, the time-independent parameters  $\xi$ , then the coefficients corresponding to the association features of the low-risk group  $\gamma_{1,m}^\ell$  and finally the ones of the high-risk group  $\gamma_{2,m}^\ell$ . The initial values of the longitudinal markers are also considered as time-independent features in these experiments. Since each longitudinal marker is associated to a vector of association features, what we call ‘‘coefficient’’ is actually the Euclidean norm of the coefficients associated to these association features, that is, the norm of  $(\gamma_{k,1}^\ell, \dots, \gamma_{k,M}^\ell)$ . In addition, to obtain standard errors for the coefficient estimates, we use a Bootstrap approach adapted to the presence of a Lasso penalty following Chzhen et al. (2019). We first run the model with the Lasso penalty to get the support of estimated coefficients. We then rerun the model 10 times without the Lasso penalty on bootstrap samples with only features whose coefficients are in the support of the first run.

Note that with this approach the errors of estimation of the coefficients that are not in the support of the first run are then not evaluated. In other words, we do not evaluate the stability of our variable selection approach. To do this, we could follow the procedure of Bach (2008) who suggest running several runs of Lasso on a bootstrap sample, and then looking at the different support of the coefficients.

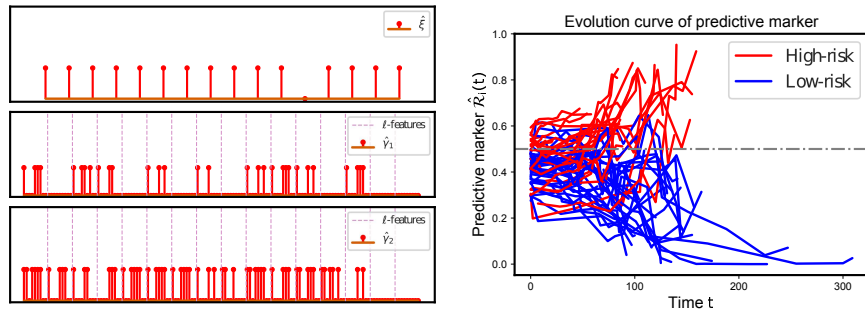
#### C.6 Experiments on a high-dimensional dataset

We evaluate the performance of FLASH model with a challenging high-dimensional dataset from NASA, which is available at <https://data.nasa.gov>. This dataset describes the degradation of 200 aircraft engines, where 17 multivariate longitudinal features are measured for each different aircraft engine until its failure. There are also three operational settings that significantly affect engine performance. Note that we only apply FLASH to this dataset because the other models did not converge after running for one day, highlighting the fact that they do not scale to high-dimensional settings.

We illustrate in Figure 8 the results obtained by FLASH. In the left panel, we can see the effect of regularization where the coefficients learned by the model are sparse and some longitudinal markers are entirely discarded. In particular, five longitudinal markers are excluded for the first group  $k = 1$  but not for the group  $k = 2$  while the last two markers are never

Table 4: Estimated coefficients with standard errors

(a) PBCseq dataset		(b) Sepsis dataset	
Features	Coefficient	Features	Coefficient
Drug	0.0	Age	$0.152 \pm 0.0414$
Age	0.0	Gender	0.0
Sex	$0.3459 \pm 0.0037$	Temp	0.0
Initial value of Serbilir	$0.3476 \pm 0.0091$	DBP	$0.0669 \pm 0.1234$
Initial value of Albumin	0.0	BaseExcess	$0.2779 \pm 0.0404$
Initial value of SGOT	$0.3664 \pm 0.0069$	HCO3	0.0
Initial value of Platelets	0.0	pH	0.0
Initial value of Prothrombin	$0.3587 \pm 0.0088$	PaCO2	$0.119 \pm 0.0987$
Initial value of Alkaline	$0.332 \pm 0.0042$	BUN	0.0
Initial value of SerChol	$0.2978 \pm 0.0093$	Calcium	0.0
Serbilir	$0.0341 \pm 0.0161$	Chloride	0.0
Albumin	$0.0833 \pm 0.0088$	Creatinine	0.0
SGOT	0.0	Glucose	0.0
Platelets	$0.0325 \pm 0.0159$	Magnesium	0.0
Prothrombin	$0.0819 \pm 0.0129$	Phosphate	0.0
Alkaline	$0.0807 \pm 0.0092$	Potassium	0.0
SerChol	0.0	Hct	0.0
Serbilir	0.0	Hgb	0.0
Albumin	0.0	PTT	0.0
SGOT	0.0	WBC	0.0
Platelets	0.0	Platelets	0.0
Prothrombin	0.0	Initial value of HR	0.0
Alkaline	$0.2788 \pm 0.0355$	Initial value of O2Sat	$0.1325 \pm 0.1049$
SerChol	0.0	Initial value of SBP	$0.6485 \pm 0.0788$
		Initial value of Resp	$0.7821 \pm 0.1001$
		HR	0.0
		O2Sat	0.0
		SBP	$0.1111 \pm 0.0211$
		Resp	$0.1016 \pm 0.0176$
		HR	$0.0198 \pm 0.0051$
		O2Sat	$0.0783 \pm 0.0066$
		SBP	$0.0849 \pm 0.0064$
		Resp	$0.0809 \pm 0.0056$



**Figure 8:** NASA dataset results. Left: in red the support of the estimated coefficient  $\hat{\xi}$  ab  $\hat{\gamma}_k$  for  $k \in \{1, 2\}$ , the dashed pink lines separate the features corresponding to each longitudinal marker  $\ell$ . Right: the evolution curves of the predictive marker  $\hat{R}_i(t)$  for varying times  $t$  and each subject where we well separate the subjects in the high risk group (in red color) and low risk (in blue color) based on their predictive marker at last measurement time with a threshold 0.5 represent by horizontal dashed line.

selected. In the right panel, we show the evolution in time of the predictive marker for each subject. We can see that, as time passes, more data is observed and the subjects are better separated into two groups of different risks.

### C.7 Experiments on using signatures as association functions

*Signature transform.* The association features can be extracted by a signature transformation. We refer to Fermanian (2021) for a detailed presentation of this transformation and simply recall here its definition. Let  $I = (\ell_1, \dots, \ell_k)$  be a word of size  $k$  from the alphabet  $\{1, \dots, L\}^k$ . The signature associated to  $I$  is defined as the mapping

$$t \mapsto \mathbf{S}^I(\mathcal{Y}_i(t^-)) := \int_{0 < u_1 < \dots < u_k < t} dy_i^{\ell_1}(u_1) \dots dy_i^{\ell_k}(u_k).$$

The signature of depth  $N$  is defined as the vector

$$\mathbf{S}_N(\mathcal{Y}_i(t^-)) = \left( \mathbf{S}^I(\mathcal{Y}_i(t^-)) \right)_{|I| \leq N}.$$

It is a transformation from a multivariate longitudinal marker to a sequence of coefficients, that are independent of time parameterization and encodes geometric properties (for example, the second order coefficients correspond to areas). It is therefore a very different transformation from the ones used in the `tsfresh` package, in particular because it encodes information on interactions between coordinates, whereas `tsfresh` focuses on univariate features. The truncation depth  $N$  is an hyperparameter that can be selected typically by cross-validation.

*Results.* Figure 9 shows the performance of FLASH with signatures as association features compared with the performance of the model with the feature extracted from `tsfresh` and the two competing methods on the four datasets presented in the article. The prediction performance of the model with the signature transformation is comparable to the competing methods and better for the FLASH\_simu dataset, and its computation cost is reduced since it does not require to implement a screening phase procedure to select the top association features.

### C.8 Procedure to select the optimal number of latent groups

The optimal number of latent groups,  $K$ , is selected based on the Bayesian information criterion (BIC) (Hastie et al., 2009), which is defined as

$$\text{BIC}(K) = -2\hat{\mathcal{L}}_n^{\text{pen}}(\theta; K) + \log(n)K,$$

where  $\hat{\mathcal{L}}_n^{\text{pen}}(\theta; K)$  is the optimal value of the likelihood function (defined in (8) the main paper) with  $K$  groups and  $n$  is the number of subjects. The optimal  $K$  is selected with the “elbow method”, that is, we pick the value of  $K$  corresponding to the first large drop in the BIC values. For example, Figure 10 shows the curve of the BIC values obtained with different  $K$  on the PBCseq dataset. In this case, the optimal  $K$  is 4.

### C.9 Sensitivity to latent class assumptions

To assess the sensitivity of the method, we have run an additional simulation where we deviate from the assumption that the latent class is generated from (1) in the main paper. More precisely, we have used a probit model (Aldrich and Nelson, 1984), where the probability of belonging to a class has the form:

$$\mathbb{P}(g_i = k) = \Phi(X_i^\top \xi_k),$$

where  $\Phi$  is the cumulative distribution function of the standard normal distribution. In Figure 11 below, we compare the performance of our method on the dataset simulated in the “well-specified” setting, described in Section 5.1 of the paper, and on a dataset simulated in this misspecified setting. We can see that the performance is slightly better in the Logistic Regression case, which was to be expected, but that this difference is not significant. Therefore, our proposed method is not too sensitive to the assumption of the model on the probability of latent class membership.

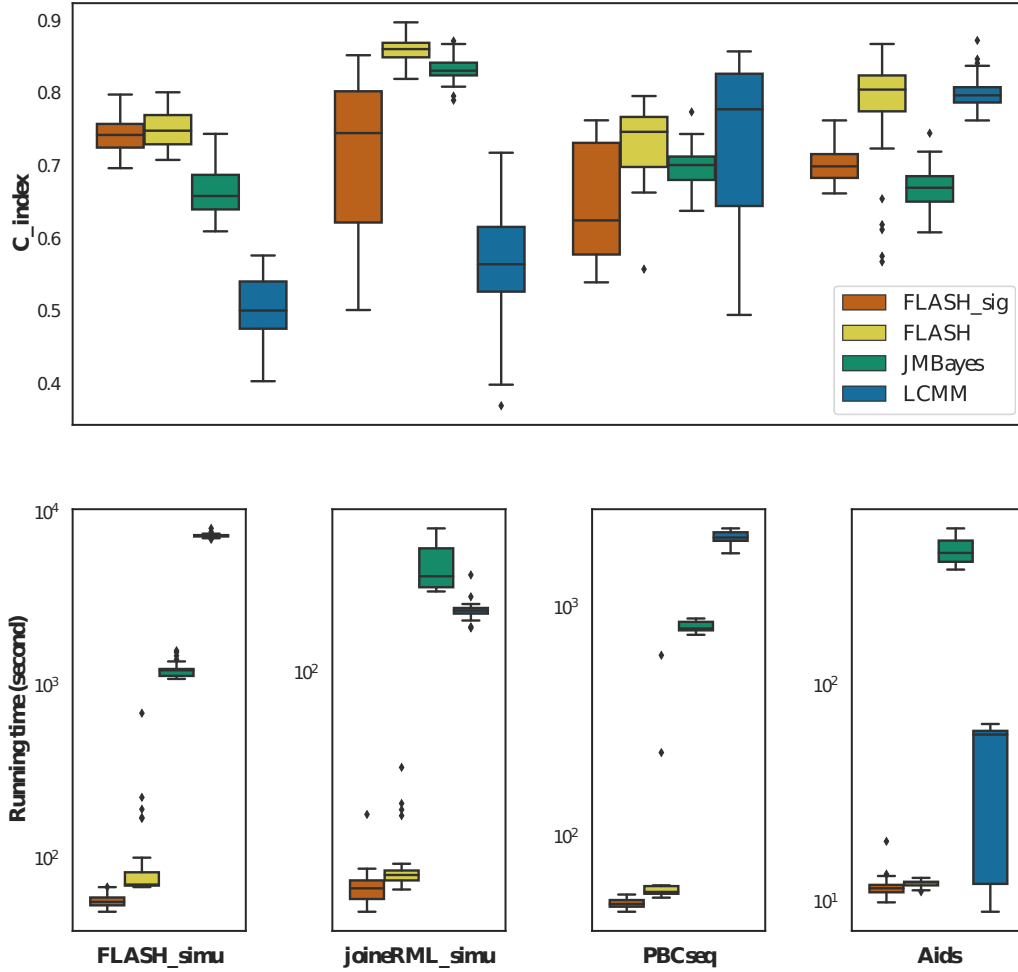


Figure 9: The performance of the Flash model with the signature transform compares with the competing methods.

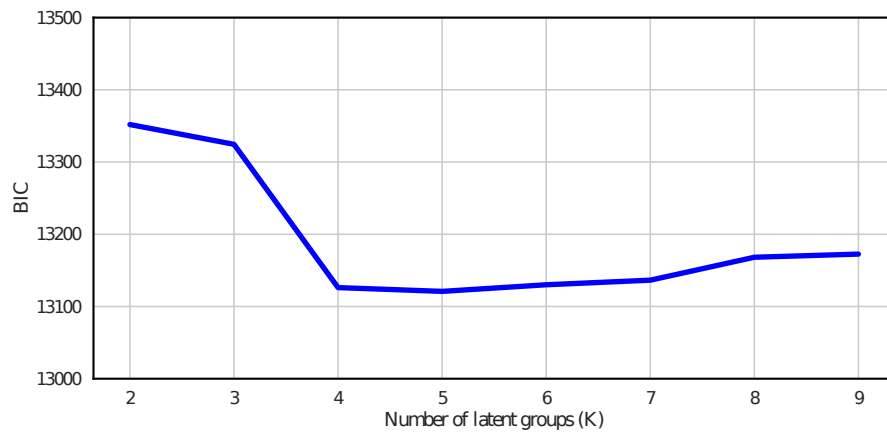
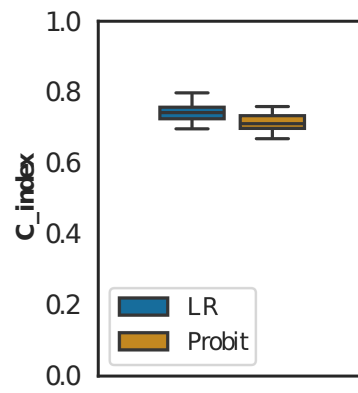


Figure 10: The BIC values on the PBCseq dataset.



**Figure 11:** The performance of FLASH in terms of C-index, on two simulated datasets: the LR (Logistic Regression) dataset is the one simulated with the model (1) in the main paper and the Probit dataset described above.