

## RESEARCH ARTICLE

# Efficient Computation of High-Dimensional Penalized Piecewise Constant Hazard Random Effects Models

Hillary M. Heiling<sup>1</sup>  | Naim U. Rashid<sup>1,2</sup> | Quefeng Li<sup>1</sup> | Xianlu L. Peng<sup>2</sup> | Jen Jen Yeh<sup>2,3,4</sup> | Joseph G. Ibrahim<sup>1</sup> 

<sup>1</sup>Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA | <sup>2</sup>Lineberger Comprehensive Cancer Center, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA | <sup>3</sup>Department of Surgery, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA | <sup>4</sup>Department of Pharmacology, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA

**Correspondence:** Hillary M. Heiling ([hillary\\_heiling@dfci.harvard.edu](mailto:hillary_heiling@dfci.harvard.edu))

**Received:** 30 January 2024 | **Revised:** 9 September 2024 | **Accepted:** 3 December 2024

**Funding:** The authors received no specific funding for this work.

**Keywords:** factor model | mixed effects | piecewise constant hazard | survival analysis | variable selection

## ABSTRACT

Identifying and characterizing relationships between treatments, exposures, or other covariates and time-to-event outcomes has great significance in a wide range of biomedical settings. In research areas such as multi-center clinical trials, recurrent events, and genetic studies, proportional hazard mixed effects models (PHMMs) are used to account for correlations observed in clusters within the data. In high dimensions, proper specification of the fixed and random effects within PHMMs is difficult and computationally complex. In this paper, we approximate the proportional hazards mixed effects model with a piecewise constant hazard mixed effects survival model. We estimate the model parameters using a modified Monte Carlo expectation conditional minimization (MCECM) algorithm, allowing us to perform variable selection on both the fixed and random effects simultaneously. We also incorporate a factor model decomposition of the random effects in order to more easily scale the variable selection method to larger dimensions. We demonstrate the utility of our method using simulations, and we apply our method to a multi-study pancreatic ductal adenocarcinoma gene expression dataset to select features important for survival.

## 1 | Introduction

Modeling survival outcomes has great clinical significance in medical and public health research. In particular, the Cox proportional hazards model has been widely utilized in order to characterize the relationship between treatments, exposures, or other covariates and patient time-to-event outcomes. However, modern biomedical datasets are increasingly high dimensional, and groups of samples within the data can exhibit complex correlations. For example, when studying survival outcomes with respect to multi-center clinical trials, recurrent events, and

genetic studies, proportional hazards mixed effects models are used to account for correlations among groups within the data and model the heterogeneity of treatment and predictor effects across groups [1, 2]. These proportional hazards mixed effects models are traditionally referred to as frailty models when the model contains a single random effect applied to the baseline hazard.

In high dimensional settings, in which the covariate effects are generally assumed to be sparse, it is often unknown a priori which covariates should be specified as fixed or random in the model.

**Abbreviations:** GLMM, generalized linear mixed model; MCECM, Monte Carlo expectation conditional minimization; PHMM, proportional hazards mixed model.

Variable selection methods such as LASSO and SCAD exist for high dimensional proportional hazards models or frailty models [3–6], but they do not allow for the selection of random effects. Several mixed effects model selection methods that rely on the specification of candidate models have been proposed, including likelihood ratios, profile Akaike information criterion (AIC) [7], and conditional AIC [8]. However, specifying all  $2^p$  possible candidate models in high dimensions is impractical. Lee et al. [9] developed a stochastic search variable selection (SSVS) method that selects both fixed and random effects in proportional hazards mixed effects models in a Bayesian framework, but their method is only computationally feasible for small or moderate dimensions.

Rashid et al. [10] developed a method that simultaneously selects both fixed and random effects in high dimensional generalized linear mixed models (GLMMs), which has since been developed into an R package available on CRAN, called **glmmPen** [11, 12]. This method broadened the feasible dimensionality of performing variable selection in GLMMs to greater dimensions than previously existing methods. This method was extended by Heiling et al. [13], who proposed a new formulation of the GLMM using a factor model decomposition of the random effects. As a result of this new formulation, they were able to improve the scalability of their method and perform variable selection within GLMMs in cases with much larger dimensions. However, these methods do not apply to survival data.

In this paper, we propose a method that simultaneously selects fixed and random effects within clustered survival data. In order to extend the methods of Rashid et al. [10] and Heiling et al. [13] to survival data, we first consider an approximation of the Cox proportional hazards model using a piecewise constant hazard model. Piecewise constant hazard models, sometimes referred to as piecewise exponential models, are parametric survival models that divide the follow-up time of the data into intervals, where the hazard function is assumed to be constant in each interval; this piecewise constant hazard survival model can be fit using a log-linear model which incorporates the duration of exposure within each interval [14–17]. Piecewise constant hazard models can be extended to include random effects using piecewise constant hazard mixed effects models [18–20]. In our piecewise constant hazard mixed effects model, we utilize the factor model decomposition of the random effects proposed in Heiling et al. [13], allowing us to scale our method to cases with hundreds of predictors. We label our method as **phmmPen\_FA**, which reflects our goal of estimating penalized proportional hazards mixed effects models using factor analysis on the random effects.

We demonstrate the application of our method by applying our method to a case study that we present in Section 4. The development of this method was motivated by a case study dataset that contains gene expression data from pancreatic ductal adenocarcinoma patients across seven separate studies. We aim to select important features that predict subjects’ survival by identifying features that increase or decrease the rate of the occurrence of death (i.e., identify features with non-zero fixed effect hazard ratios) as well as identify features that have varied effects on subjects’ survival across the groups (i.e., identify non-zero random effects). Due to the large number of features that we consider, it is difficult to have a priori knowledge of which features

have non-zero fixed or random effects. Therefore, we will use the **phmmPen\_FA** method to fit a penalized piecewise constant hazard mixed effects survival model to select important fixed and random effects.

The remainder of this paper is organized as follows. Section 2 reviews the statistical models and algorithm used to estimate piecewise constant hazard mixed effects models. In Section 3, simulations are conducted to assess the performance of our method. Section 4 describes the motivating case study for the prediction of survival in pancreatic ductal adenocarcinoma cancer using gene expression data from multiple studies, and provides results from the application of our new method to the case study. We close the article with some discussion in Section 5.

Software in the form of R code for the **phmmPen\_FA** procedure is available through the **glmmPen** package in CRAN <https://cran.r-project.org/package=glmmPen> and the GitHub repository <https://github.com/hheiling/glmmPen>. The **phmmPen\_FA** procedure is implemented through the `phmmPen_FA()` function within this **glmmPen** R package.

## 2 | Methods

### 2.1 | Model Formulation

In this section, we review the notation and model formulation of our approach. We begin by introducing the proportional hazards mixed effects model using a generic linear predictor with both fixed and random effects. We then discuss how we approximate this model with a piecewise constant hazard mixed effects model, which is similar to how fixed-effects only Cox proportional hazards models are approximated with piecewise constant hazard models [14, 19–21]. Next, we explain how we represent this model using a log-linear mixed effects model [15, 16, 19]. Once we have set up the general model scheme, we provide details about the exact form of our assumed linear predictor, where we incorporate the factor model decomposition of the random effects proposed in Heiling et al. [13] Lastly, we introduce penalties into the model and discuss how the variable selection goals of our algorithm relate to the introduced notation. Details about the MCECM algorithm we use to fit this model are discussed in Section 2.2.

We consider the case where we want to analyze data from  $K$  independent groups of subjects. For each group  $k = 1, \dots, K$ , there are  $n_k$  subjects for a total sample size of  $N = \sum_{k=1}^K n_k$ . For group  $k$ , let  $\mathbf{y}_k = (y_{k1}, \dots, y_{kn_k})^T$  be the vector of  $n_k$  observed times, where  $y_{ki} = \min(T_{ki}, C_{ki})$ ,  $T_{ki}$  represents the event time, and  $C_{ki}$  represent censoring time; let  $\delta_k = (\delta_{k1}, \dots, \delta_{kn_k})$  where  $\delta_{ki} = I(T_{ki} < C_{ki})$  represents the indicator that a subject’s event time was observed; and let  $\mathbf{x}_{ki} = (x_{ki,1}, \dots, x_{ki,p})^T$  be the  $p$ -dimensional vector of predictors, and  $\mathbf{X}_k = (\mathbf{x}_{k1}, \dots, \mathbf{x}_{kn_k})^T$ . We standardize the fixed effects covariates matrix  $\mathbf{X} = (\mathbf{X}_1^T, \dots, \mathbf{X}_K^T)^T$  such that  $\sum_{k=1}^K \sum_{i=1}^{n_k} x_{ki,j} = 0$  and  $N^{-1} \sum_{k=1}^K \sum_{i=1}^{n_k} x_{ki,j}^2 = 1$  for  $j = 1, \dots, p$ .

We would like to estimate the proportional hazards mixed effects model

$$h(t|\eta_{ki}) = h_0(t) \exp(\eta_{ki}), \quad (1)$$

where  $h(t|\eta_{ki})$  is the individual hazard of subject  $i$  in group  $k$  at time  $t$ ,  $h_0(t)$  represents the baseline hazard at time  $t$ , and  $\eta_{ki}$  represents the linear predictor containing the fixed effects log hazard ratio coefficients, the group-specific random effects, and the subject's individual covariates. The exact form of the linear predictor  $\eta_{ki}$  assumed in our model is described later in this section.

In cases where we are modeling survival data with only fixed effects in the model (i.e., no random effects) and we are not using Bayesian techniques to estimate these fixed effects, we can ignore the baseline hazard function  $h_0(t)$  during the estimation of the fixed effects coefficients in the linear predictor. However, when survival models include random effects or involve estimation using Bayesian techniques, it is necessary to fully specify and model the baseline hazard function. In this paper, we approximate the baseline hazard using a piecewise constant function [14] such that we approximate the model (1) using the piecewise constant hazard mixed effects model [19]. We use this approximation because modeling the baseline hazard as a piecewise constant function in this setting allows for relatively convenient computation. We first partition the time of the study into  $J$  intervals, where we assume that the baseline hazard within a particular time interval is constant. Please see Section 3 for a discussion on choosing the number of time intervals  $J$ . Let us define the cut points  $0 = \tau_0 < \tau_1 < \dots < \tau_J = \infty$ , and let  $h_j$  be the constant baseline hazard within interval  $j$ ,  $[\tau_{j-1}, \tau_j)$ . We then approximate (1) using the model

$$h_{kij} = h_j \exp(\eta_{ki}), \quad (2)$$

where  $h_j$  is the baseline hazard for interval  $j$  and  $h_{kij}$  is the constant hazard corresponding to subject  $i$  in group  $k$  within interval  $j$ .

The observed data for each subject includes their observed time  $y_{ki}$  and their event indicator  $\delta_{ki}$ . We extend these to define analogous measures for each interval, where  $t_{kij}^* = \max[\min(y_{ki}, \tau_j) - \tau_{j-1}, 0]$  is the amount of time subject  $i$  in group  $k$  survived within interval  $j$ , and  $d_{kij} = I(\tau_{j-1} \leq y_{ki} < \tau_j, \delta_{ki} = 1)$  is the indicator of whether the subject died during interval  $j$ . To better clarify  $t_{kij}^*$ , this term has three possible values, determined by the relative value of their observed time  $y_{ki}$  to the interval cut points:

$$t_{kij}^* = \begin{cases} \tau_j - \tau_{j-1}, & y_{ki} > \tau_j; \\ y_{ki} - \tau_{j-1}, & \tau_{j-1} < y_{ki} \leq \tau_j; \\ 0, & y_{ki} \leq \tau_{j-1}. \end{cases}$$

We can then treat the death indicators  $d_{kij}$  as if they are independent Poisson observations with means  $\mu_{kij} = t_{kij}^* h_{kij}$ . Consequently, we estimate the fixed and random effect parameters of our model by fitting the data using the log-linear model

$$\log \mu_{kij} = \log t_{kij}^* + \psi_j + \eta_{ki}, \quad (3)$$

where  $\psi_j = \log(h_j)$  is the logarithm of the constant hazard within interval  $j$  and  $\log(t_{kij}^*)$ , which represents the log of the time a subject survived within interval  $j$ , is treated as an offset to the model. We treat the  $d_{kij}$  values as the outcome (i.e., response) values of model (3) for each subject  $i$  in group  $k$  for each interval  $j$ .

Let us define  $\mathbf{d}_k = (d_{k11}, \dots, d_{k1J}, \dots, d_{kn_k1}, \dots, d_{kn_kJ})^T$  as the vector of death indicator values for all subjects in group  $k$  and all  $J$  time intervals. Then, the piecewise constant hazard likelihood is defined as

$$f(\mathbf{d}_k | \mathbf{X}_k, \alpha_k; \theta) = \prod_{i=1}^{n_k} \prod_{j=1}^J [I(t_{kij}^* > 0) \mu_{kij}]^{d_{kij}} \exp[-I(t_{kij}^* > 0) \mu_{kij}], \quad (4)$$

where  $\mu_{kij}$  is defined in (3), and  $I(t_{kij}^* > 0) = 1$  indicates that a subject  $i$  in group  $k$  survived at least part way through interval  $j$ , 0 if the subject died or was censored before interval  $j$ .

Now we may define the form of the linear predictor term  $\eta_{ki}$  used within this model. We start by introducing the linear predictor term used in the traditional generalized linear mixed model [10, 22, 23]

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\gamma}_k = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\Gamma} \boldsymbol{\epsilon}_k, \quad (5)$$

where  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$  is a  $p$ -dimensional vector for the fixed effects coefficients ( $\boldsymbol{\beta}$  represents the log hazard ratio values for each predictor and excludes an intercept, just as in a typical Cox proportional hazards model),  $\boldsymbol{\Gamma}$  is the Cholesky decomposition of the random effects covariance matrix  $\boldsymbol{\Sigma}$  such that  $\boldsymbol{\Gamma} \boldsymbol{\Gamma}^T = \boldsymbol{\Sigma}$ ,  $\boldsymbol{\gamma}_k = \boldsymbol{\Gamma} \boldsymbol{\epsilon}_k$  is a  $q$ -dimensional vector of unobservable random effects (including the random intercept) for group  $k$  where  $\boldsymbol{\epsilon}_k \sim N_q(0, \mathbf{I})$ , and  $\mathbf{z}_{ki}$  is a  $q$ -dimensional vector that includes an intercept term and a subset of  $\mathbf{x}_{ki}$ .

We reformulate the linear predictor of (5) using the technique described in Heiling et al. [13], where we decompose the random effects  $\boldsymbol{\gamma}_k$  into a factor model with  $r$  latent common factors and we assume  $r \ll q$ . As a result, we assume  $\boldsymbol{\gamma}_k = \mathbf{B} \boldsymbol{\alpha}_k$ , where  $\mathbf{B}$  is the  $q \times r$  loading matrix and  $\boldsymbol{\alpha}_k$  represents the  $r$  latent common factors. We assume the latent factors  $\boldsymbol{\alpha}_k$  are uncorrelated and follow a  $N_r(\mathbf{0}, \mathbf{I})$  distribution. We re-write the linear predictor as

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k. \quad (6)$$

In the representation of (6), the random component of the linear predictor has variance  $\text{Var}(\mathbf{B} \boldsymbol{\alpha}_k) = \mathbf{B} \mathbf{B}^T = \boldsymbol{\Sigma}$ , which is low rank. By using this representation, we reduce the dimension of the latent space from  $q$  to  $r$ . As a result, this reduces the dimension of the integral in the likelihood, which reduces the computational complexity of the E-step in the EM algorithm described in Section 2.2. Consequently, this factor decomposition reduces the computational time of the algorithm and enables our method to scale to hundreds of predictors [13].

In order to estimate  $\mathbf{B}$ , let  $\mathbf{b}_t \in \mathbb{R}^r$  be the  $t$ -th row of  $\mathbf{B}$  and  $\mathbf{b} = (\mathbf{b}_1^T, \dots, \mathbf{b}_q^T)^T$ . We then reparameterize the linear predictor as

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k = \left( \mathbf{x}_{ki}^T \quad (\boldsymbol{\alpha}_k \otimes \mathbf{z}_{ki})^T \mathbf{J}^\circ \right) \begin{pmatrix} \boldsymbol{\beta} \\ \mathbf{b} \end{pmatrix} \quad (7)$$

in a manner similar to Chen and Dunson [22] and Ibrahim et al. [23], where  $\mathbf{J}^\circ$  is a matrix that transforms  $\mathbf{b}$  to  $\text{vec}(\mathbf{B})$  such that  $\text{vec}(\mathbf{B}) = \mathbf{J}^\circ \mathbf{b}$  and  $\mathbf{J}^\circ$  is of dimension  $(qr) \times (qr)$ . The vector of parameters  $\theta = (\beta^T, \mathbf{b}^T, \boldsymbol{\psi}^T)^T$  are the main parameters of interest.

We denote the true value of  $\theta$  as  $\theta^* = (\beta^{*T}, \mathbf{b}^{*T}, \boldsymbol{\psi}^{*T})^T = \text{argmin}_\theta E_\theta[-\ell(\theta)]$  where  $\ell(\theta)$  is the observed log-likelihood across all  $K$  groups such that  $\ell(\theta) = \sum_{k=1}^K \ell_k(\theta)$ , where  $\ell_k(\theta) = (1/n_k) \log \int f(\mathbf{d}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \theta) \phi(\boldsymbol{\alpha}_k) d\boldsymbol{\alpha}_k$ .

Our primary goal is to select the true nonzero fixed and random effects, that is, identify the set

$$S = S_1 \cup S_2 = \{j : \beta_j^* \neq 0\} \cup \{t : \|\mathbf{b}_t^*\|_2 \neq 0\},$$

where  $S_1$  and  $S_2$  represent the true fixed and random effects, respectively. When  $\mathbf{b}_t = \mathbf{0}$ , this indicates that the effect of covariate  $t$  is fixed across the  $K$  groups (i.e., the corresponding  $t$ -th row and column of  $\boldsymbol{\Sigma}$  is set to 0).

Our objective is to solve the penalized likelihood problem of (8):

$$\hat{\theta} = \text{argmin}_\theta -\ell(\theta) + \lambda_0 \sum_{l=1}^p \rho_0(\beta_l) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2), \quad (8)$$

where  $\ell(\theta)$  is the observed log-likelihood for all  $K$  groups,  $\rho_0(t)$  and  $\rho_1(t)$  are general folded-concave penalty functions, and  $\lambda_0$  and  $\lambda_1$  are positive tuning parameters. The penalty functions applied to the fixed effects, represented by  $\rho_0(t)$ , could include the  $L_1$  penalty (LASSO), the smoothly clipped absolute deviation (SCAD) penalty, and the minimax concave penalty (MCP) penalty [24, 25]. The penalty functions applied to the random effects, represented by  $\rho_1(t)$ , could include the group LASSO, the group MCP, or the group SCAD penalties presented by Breheny and Huang [26] since we treat the elements of  $\mathbf{b}_t$  as a group and penalize them in a group wise manner. As a result, these groups of  $\mathbf{b}_t$  are estimated to be either all zero or all nonzero, which means that we select covariates to have random effects ( $\hat{\mathbf{b}}_t \neq \mathbf{0}$ ) or fixed effects ( $\hat{\mathbf{b}}_t = \mathbf{0}$ ) across the  $K$  groups.

Before moving on to the specifics of the algorithm used to perform this variable selection procedure and estimate the fixed and random effect parameters of interest, we want to provide some further clarifications about the values of  $p$ ,  $q$ , and  $r$  discussed in both this section and remaining sections. The values of  $p$  and  $q$  refer to the full set of fixed and random effect predictors, respectively, that are considered in the model and would be input into the `phmmPen_FA` variable selection procedure. Let us represent the true number of non-zero fixed and random effects of in the model as  $p^* \leq p$  and  $q^* \leq q$ , respectively. When we claim that  $r \ll q$  (i.e.,  $r$  is much less than  $q$ ), this refers to the relative size of the number of latent factors used in the model versus the total number of random effects considered in the model. We do not necessarily assume  $r \ll q^*$  holds for the true number of random effects in the model, although we do assume  $r < q^*$ .

## 2.2 | MCECM Algorithm

We solve (8) for some specific  $(\lambda_0, \lambda_1)$  using a Monte Carlo expectation conditional minimization (MCECM) algorithm [27].

Our objective within the  $s$ th iteration of the MCECM algorithm is to evaluate the expectation of (E-step) and minimize (M-step) the penalized Q-function defined in (9):

$$\begin{aligned} Q_\lambda(\theta | \theta^{(s)}) &= \sum_{k=1}^K E \left\{ -\log(f(\mathbf{d}_k, \mathbf{X}_k, \boldsymbol{\alpha}_k; \theta | \mathbf{D}_o; \theta^{(s)})) \right\} \\ &+ \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2) \\ &= Q_1(\theta | \theta^{(s)}) + Q_2(\theta^{(s)}) \\ &+ \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2), \end{aligned} \quad (9)$$

where  $(\mathbf{d}_k, \mathbf{X}_k, \boldsymbol{\alpha}_k)$  gives the complete data for group  $k$ ,  $\mathbf{D}_{k,o} = (\mathbf{d}_k, \mathbf{X}_k)$  gives the observed data for group  $k$ ,  $\mathbf{D}_o$  represents the entirety of the observed data, and

$$Q_1(\theta | \theta^{(s)}) = - \sum_{k=1}^K \int \log [f(\mathbf{d}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \theta)] \phi(\boldsymbol{\alpha}_k | \mathbf{D}_{k,o}; \theta^{(s)}) d\boldsymbol{\alpha}_k, \quad (10)$$

$$Q_2(\theta^{(s)}) = - \sum_{k=1}^K \int \log [\phi(\boldsymbol{\alpha}_k)] \phi(\boldsymbol{\alpha}_k | \mathbf{D}_{k,o}; \theta^{(s)}) d\boldsymbol{\alpha}_k \quad (11)$$

where  $f(\mathbf{d}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \theta)$  was defined in (4) and  $\phi(\boldsymbol{\alpha}_k | \mathbf{D}_{k,o}; \theta^{(s)})$  represents the posterior distribution of the latent factors  $\boldsymbol{\alpha}_k$ .

Our goal in the E-step of the algorithm is to approximate the  $r$ -dimensional integral expressed in (10). We first specify  $J$  time intervals (see Section 3 for a discussion on choosing a value of  $J$ ) defined so that there are an approximately equal number of events within each time interval [14]. If a subject survived at least part-way through  $j^*$  intervals (i.e.,  $t_{kij}^* > 0$  for  $j = 1, \dots, j^* \leq J$ ), the long-form dataset contains  $j^*$  observations for that subject. For subject  $i$  in group  $k$  that survived at least part-way through  $j^*$  time intervals, we define  $d_{kij} = I(\tau_{j-1} \leq y_{ki} < \tau_j, \delta_{ki} = 1)$  for  $j = 1, \dots, j^* \leq J$ , the subject's  $\mathbf{x}_{ki}$  and  $\mathbf{z}_{ki}$  covariates are repeated for all  $j^*$  observations, the  $\log(t_{kij}^*)$  offset term is calculated for each interval, and additional reference coded indicator values  $\mathbf{v}_{kij} = (v_{kij,1}, \dots, v_{kij,J})^T$  for the time interval  $j = 1, \dots, j^*$  are specified. The first element  $\mathbf{v}_{kij}$  (element  $v_{kij,1}$ ) is always 1, encoding a fixed effect intercept which represents time interval 1. For time interval  $j > 1$ , the  $j$ -th element  $\mathbf{v}_{kij}$  (element  $v_{kij,j}$ ) is also 1. All other elements of  $\mathbf{v}_{kij}$  are set to 0.

Instead of estimating  $\boldsymbol{\psi}$  directly, we reformulate this quantity as  $\tilde{\boldsymbol{\psi}}$ , where  $\psi_1 = \tilde{\psi}_1$  and  $\psi_j = \tilde{\psi}_1 + \tilde{\psi}_j$  for  $j \geq 2$ . In this formulation,  $\tilde{\psi}_1$  estimates the log of the baseline hazard for time interval  $[\tau_0, \tau_1)$ , and  $\tilde{\psi}_1 + \tilde{\psi}_j$  estimates the log of the baseline hazard for time interval  $[\tau_{j-1}, \tau_j)$  for  $j = 2, \dots, J$ . By estimating the log baseline hazard parameters in this way, we are including a fixed effect intercept in our model. By including a fixed effect intercept, we ensure that the full  $\mathbf{z}_{ki}$  vector, which includes a random intercept, is a subset of the subject's fixed effects.

We can re-write the log-linear model of (3) as

$$\log \mu_{kij} = \log t_{kij}^* + \mathbf{v}_{kij}^T \tilde{\boldsymbol{\psi}} + \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k. \quad (12)$$

### 2.2.1 | Monte-Carlo E-Step

The integrals in the Q-function do not have closed forms. We approximate these integrals using a Markov Chain Monte Carlo (MCMC) sample of size  $M$  from the posterior density  $\phi(\boldsymbol{\alpha}_k | \mathbf{D}_{k,o}; \boldsymbol{\theta}^{(s)})$ . Let  $\boldsymbol{\alpha}_k^{(s,m)}$  be the  $m^{\text{th}}$  simulated  $r$ -dimensional vector from the posterior of the latent common factors,  $m = 1, \dots, M$ , at the  $s^{\text{th}}$  iteration of the algorithm for group  $k$ . The integral in (10) can be approximated as:

$$\begin{aligned} Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) &\approx -\frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K \log f(\mathbf{d}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k^{(s,m)}; \boldsymbol{\theta}) \\ &= -\frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K \sum_{i=1}^{n_k} \sum_{j=1}^J I(t_{kij}^* > 0) [d_{kij} \log \mu_{kij}^{(s,m)} - \mu_{kij}^{(s,m)}], \end{aligned} \quad (13)$$

where  $\log \mu_{kij}^{(s,m)} = \log t_{kij}^* + \mathbf{v}_{kij}^T \tilde{\boldsymbol{\psi}} + \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k^{(s,m)}$ . We use the No-U-Turn Sampler Hamiltonian Monte Carlo sampling procedure (NUTS HMC) from the Stan software [28, 29] to sample the latent factors  $\boldsymbol{\alpha}_k^{(s,m)}$ ; we use this sampling procedure (implemented using the **rstan** R package [30]) because this helps improve the overall speed and efficiency of our E-step compared to other appropriate sampling method options, thereby helping to improve the speed of the overall MCECM algorithm. While we used the NUTS HMC sampling procedure for all of the analyses described in this paper, our software (see the Data S1) allows for other MCMC sampling procedures including the Metropolis-within-Gibbs with an adaptive random walk sampler [31] and the Metropolis-within-Gibbs with an independence sampler [32].

### 2.2.2 | M-Step

In the M-step of the algorithm, we aim to minimize

$$Q_{1,\lambda}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) = Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) + \lambda_0 \sum_{l=1}^p \rho_0(\beta_l) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2) \quad (14)$$

with respect to  $\boldsymbol{\theta} = (\boldsymbol{\beta}^T, \mathbf{b}^T, \tilde{\boldsymbol{\psi}}^T)^T$ . We do this by using a Majorization-Minimization algorithm with penalties applied to the fixed effects  $\boldsymbol{\beta}$  and the rows of  $\mathbf{B}$ . The step size of the Majorization-Minimization algorithm is estimated using a proximal gradient line search algorithm [33].

Let  $s$  represent the iteration of the MCECM algorithm, and let  $g$  represent the iteration within a particular M-step of the MCECM algorithm such that the coefficients for the  $g$ -th M-step update within the  $s$ -th MCECM iteration is  $\boldsymbol{\theta}^{(s,g)}$ .

**Initialization:** We initialize the parameters  $\boldsymbol{\theta}^{(s,0)}$  using  $\boldsymbol{\theta}^{(s-1)}$ . The step size for the Majorization-Minimization algorithm,  $c^{(s,0)}$ , is initialized using  $c^{(s-1)}$  (value of 1.0 if  $s = 1$ ). Details on the initialization of  $\boldsymbol{\theta}^{(0)}$  is given in Section 2.2.3.

**Coefficient updates:** Conditional on  $\boldsymbol{\beta}^{(s,g-1)}$  and  $\mathbf{b}^{(s,g-1)}$ , each  $\tilde{\boldsymbol{\psi}}_j^{(s,g)}$  for  $j = 1, \dots, J$  is given a single update using the

Majorization-Minimization algorithm specified by Breheny and Huang [26] with no penalization applied.

Conditional on  $\mathbf{b}^{(s,g-1)}$  and the recently updated  $\tilde{\boldsymbol{\psi}}^{(s,g)}$ , each  $\beta_l^{(s,g)}$  for  $l = 1, \dots, p$  is given a single update using the Majorization-Minimization algorithm specified by Breheny and Huang [26].

Conditional on the recently updated  $\boldsymbol{\beta}^{(s,g)}$  and  $\tilde{\boldsymbol{\psi}}^{(s,g)}$ , each  $\mathbf{b}_t^{(s,g)}$  for  $t = 1, \dots, q$  is updated using the Majorization-Minimization coordinate descent grouped variable selection algorithm specified by Breheny and Huang [26].

If necessary, the step size  $c^{s,g+1}$  is updated using a proximal gradient line search algorithm [33] and multiplied by a factor of 0.95.

**Convergence:** The coefficient update steps are repeated until the convergence criteria specified in Section S1.4 is reached or until the M-step reaches its maximum number of iterations (default 50).

### 2.2.3 | MCECM Algorithm

The full MCECM algorithm for estimating the parameters with a particular  $(\lambda_0, \lambda_1)$  proceeds as described below.

**Initialization:** Very briefly, we initialize  $\boldsymbol{\theta}^{(0)}$  using either the coefficients from a previous model fit or from a naive model assuming no random effects if no previous model fit is available. We include additional comments on initialization later in this section; see Section S1.4 for full details.

**E-step:** For EM iteration  $s$ , a burn-in sample from the posterior distribution of the latent factors is run and discarded. A sample of size  $M^{(s)}$  from the posterior  $(\tilde{\boldsymbol{\alpha}}_k^{(s)} = \left( \left( \boldsymbol{\alpha}_k^{(s,1)} \right)^T, \dots, \left( \boldsymbol{\alpha}_k^{(s,M)} \right)^T \right)^T$  for  $k = 1, \dots, K$ ) is then drawn and retained for the M-step.

**M-step:** Parameter estimates of  $\tilde{\boldsymbol{\psi}}^{(s)}$ ,  $\boldsymbol{\beta}^{(s)}$ , and  $\mathbf{b}^{(s)}$  are then updated as described in the M-step procedure given in Section 2.2.2.

**Convergence:** The E-step and M-step are repeated until the MCECM convergence condition specified in Section S1.4 is met two consecutive times (default) or until the maximum number of EM iterations is reached (25 in our simulations).

Sections S1.2 and S1.3 outline the process of model selection and finding optimal tuning parameters. In brief, the algorithm runs a computationally efficient two-stage approach to pick the optimal set of tuning parameters. In the first stage of this approach, the algorithm keeps the fixed effect penalty constant at the minimum value of the fixed effects penalty sequence  $\lambda_{0,\min}$ , and searches over the sequence of the random effects penalties from the minimum random effect penalty  $\lambda_{1,\min}$  to the maximum value  $\lambda_{1,\max}$ . We use model selection criteria (BIC-ICQ [23]) to select the best random effect penalty, or  $\lambda_{1,\text{opt}}$ . In the second stage, the algorithm keeps the random effect penalty fixed at  $\lambda_{1,\text{opt}}$  and searches over the sequence of fixed effect penalties from  $\lambda_{0,\min}$  to  $\lambda_{0,\max}$ . The overall best model is chosen from the models in the second stage

using model selection criteria. Details on the calculation of the  $\lambda_0$  and  $\lambda_1$  sequences are given in Section S1.3.

During the development of this method, we found that certain initialization procedures helped improve variable selection and estimation results. Therefore, we briefly summarize our initialization procedure here, with full details of our initialization and convergence procedures provided in Section S1.4. First, we outline how fixed and random effect coefficients are initialized in the first model (using minimum penalties  $\lambda_{0,\min}$  and  $\lambda_{1,\min}$ ) in the sequence of models fit for the variable selection procedure.

We initialize the fixed effects coefficients (baseline constant hazard coefficient values  $\tilde{\psi}^{(0)}$  and predictor fixed effect coefficient values  $\beta^{(0)}$ ) by first fitting a penalized piecewise constant hazard model assuming only fixed effects and no random effects. In this model, the  $\tilde{\psi}^{(0)}$  coefficients are not penalized, and we penalize the  $\beta^{(0)}$  coefficients using the minimum fixed effects penalty  $\lambda_{0,\min}$ . The coefficient values from this minimum penalty model assuming no random effects are used as input for the first model of the overall variable selection procedure.

Based on the initialized  $\beta^{(0)}$ , the predictors with non-zero initialized fixed effects are also initialized to have non-zero random effects (i.e., the corresponding rows of the  $\mathbf{B}^{(0)}$  matrix are set to non-zero values), and predictors with zero-valued initialized fixed effects are initialized to have zero-valued random effects (i.e., the corresponding rows of the  $\mathbf{B}^{(0)}$  matrix are set to zero). This initial screening of random effects helps improve the speed of the algorithm. See Section S1.4 for details on how the non-zero elements of  $\mathbf{B}^{(0)}$  are initialized.

After we fit the first model in the overall variable selection procedure using the MCECM algorithm, the fixed and random effect coefficients in consecutive models are initialized using the values from the previous model fit. We found that progressing through penalty values from the minimum penalty to the maximum penalty (as discussed above and in Sections S1.2 and S1.3) significantly improved initialization of subsequent models, as opposed to proceeding from the maximum penalty to the minimum penalty as some other fixed-effects only penalization methods do [24, 25].

### 2.3 | Estimation of the Number of Latent Factors

Performing our proposed **phmmPen\_FA** method requires specifying the number of latent factors  $r$ . Since  $r$  is typically unknown a priori, this value needs to be estimated. Here, we use the Growth Ratio (GR) procedure [34].

The GR method for our application requires a  $q \times K$  matrix of observed random effects. Since these random effects cannot be directly observed, we instead calculate pseudo random effects by first fitting a penalized piecewise constant survival model with a small penalty to each group individually using the random effect covariates of interest as the predictors in the model. We then take these group-specific estimates and center them so that all features have a mean of 0. Let these  $q$ -dimensional group-specific estimates be denoted as  $\hat{\gamma}_k$  for each group  $k = 1, \dots, K$ . We then

define  $\mathbf{G} = (\hat{\gamma}_1, \dots, \hat{\gamma}_K)$  as the final  $q \times K$  matrix of pseudo random effects.

Let  $\xi_j(A)$  be the  $j$ -th largest eigenvalue of the positive semidefinite matrix  $A$ , and let  $\tilde{\mu}_{qK,j} \equiv \xi_j(\mathbf{G}\mathbf{G}^T/(qK)) = \xi_j(\mathbf{G}^T\mathbf{G}/(qK))$ . To find the GR estimator, we first order the eigenvalues of  $\mathbf{G}\mathbf{G}^T/(qK)$  from largest to smallest. Then, we calculate the following ratios:

$$GR(j) \equiv \frac{\log[V(j-1)/V(j)]}{\log[V(j)/V(j+1)]} = \frac{\log(1 + \tilde{\mu}_{qK,j}^*)}{\log(1 + \tilde{\mu}_{qK,j+1}^*)}, \quad j = 1, 2, \dots, U \quad (15)$$

where  $V(j) = \sum_{l=j+1}^{\min(q,K)} \tilde{\mu}_{qK,l}$ ,  $\tilde{\mu}_{qK,j}^* = \tilde{\mu}_{qK,j}/V(j)$ , and  $U$  is a pre-defined constant. Then, we estimate  $r$  by

$$\hat{r}_{GR} = \max_{1 \leq j \leq U} GR(j) \quad (16)$$

## 3 | Simulations

In this section, we examine how well the **phmmPen\_FA** algorithm performs variable selection on the fixed and random effects covariates for piecewise constant hazard mixed effects models under several different conditions. In all of these simulations, we use the MCP penalty (MCP penalty for the fixed effects, group MCP penalty for the rows of the  $\mathbf{B}$  matrix) and the BIC-ICQ [23] model selection criterion with the abbreviated two-stage grid search as described in the Section 2.2 (see full details in Section S1.2). In order to determine the robustness of our variable selection procedure based on the assumed value of  $r$ , we fit models in one of two ways: we estimated the number of common factors  $r$  using the Growth Ratio estimation procedure, or we use the true value of  $r$ .

In all of our simulations, we specified  $J = 8$  time intervals for the piecewise constant hazard mixed effect survival model; we chose  $J = 8$  because this is the default number of time intervals used within the piecewise constant hazard procedure implemented in the SAS **proc phreg** command [35]. We created the intervals such that there were an approximately equal number of events within each time interval, as suggested by Allison [14]. There are several alternative options for choosing the number of time intervals to include in a piecewise constant hazard mixed effect model. Castillo and van der Pas [36] proposed using  $J = \sqrt{n/\log n}$  as a data-driven estimate for the number of intervals. Our review of the literature suggests that it is common to choose between 5 and 10 time intervals for piecewise constant hazard models when they are fitting real data [17, 19, 35, 37, 38]. One could try several values of  $J$  and compare the models using appropriate Bayesian model selection criteria, including the DIC [39–41], the  $L$  measure [18, 42], or the BIC-ICQ [23]. In general, it is suggested to use a moderate number of time intervals because using too many time intervals can result in unstable estimates, but too few time intervals can lead to an inadequate model fit [38]. In the Section S3.5, we repeat some simulations assuming a range between 5 and 10 time intervals in our **phmmPen\_FA** procedure. The variable selection performance was very similar across this range of time intervals, see the [Supporting Information S2](#) for full details. Our software in the package **phmmPen\_FA** recommends choosing between 5 and 10 time intervals and specifies a default of eight time intervals.

### 3.1 | Variable Selection in Survival Data With 100 Predictors

We examine the performance and scalability of the **phmm-Pen\_FA** algorithm when performing variable selection in high dimensions of  $p = 100$  total predictors. We simulated survival data from a piecewise constant hazard mixed effect model with  $p$  predictors. Of  $p$  total predictors, we assume that the first five predictors have truly non-zero fixed and random effects (i.e.,  $p^* = 5$  and  $q^* = p^* + 1$ , where the additional 1 comes from including a random intercept to account for group-specific variations in the baseline hazard), and the other  $p - 5$  predictors have zero-valued fixed and random effects. We specified a full model for the algorithm such that the random effect predictors equaled the fixed effect predictors (e.g., in the full model  $q = p + 1$ ), and our aim was to select the set of true predictors and random effects.

To simulate the data, we set the total sample size to  $N = 1000$  and considered the number of groups  $K$  to be either 5 or 10, with an equal number of subjects per group. We set up the random effects covariance matrix by specifying a  $\mathbf{B}$  matrix with dimensions  $(p + 1) \times r$ , where  $p + 1$  represents the  $p$  predictors specified in the  $\mathbf{X}$  matrix plus the random intercept, and the number of latent common factors  $r$  was set to three. Six of these  $p + 1$  rows—corresponding to the true five predictors plus the random intercept—had non-zero elements, while the remaining  $p - 5$  rows were set to zero. For each value of  $r$ , we considered a  $\mathbf{B}$  matrix that produced  $\mathbf{\Sigma} = \mathbf{B}\mathbf{B}^T$  with either small or moderate variances and eigenvalues; see Section S1.1 for further details. These two cases are referred to as the ‘small’ or ‘moderate’  $\mathbf{B}$  matrices in the simulation results presented in this section. We generate both moderate and strong predictor effects, where all 5 of the true fixed effects have coefficient values of 0.5 or 1.0, respectively. Each condition was evaluated using 100 total simulated datasets.

In order to sample event times  $\mathbf{T} = (\mathbf{T}_1^T, \dots, \mathbf{T}_k^T)^T$  where  $\mathbf{T}_k = (T_{k1}, \dots, T_{kn_k})^T$ , we defined five half-year time intervals as  $\{[0,0.5), [0.5,1.0), [1.0,1.5), [1.5,2.0), [2.0, \infty)\}$ . The corresponding log baseline hazard values for these intervals were  $\psi_j^* = (-1.5, 1.0, 2.7, 3.7, 6.8)$ .

For group  $k$ , we generated the event times  $T_{ki}$ ,  $i = 1, \dots, n_k$ , using the following procedure: We first simulated values from the exponential distribution  $e_{kij} \sim \text{Exp}(R_{kij})$  starting with  $j = 1$ , where the exponential rate  $R_{kij} = \exp(\psi_j + \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\gamma}_k)$ , where  $\boldsymbol{\gamma}_k \sim N_6(0, \mathbf{B}\mathbf{B}^T)$ . If the inequality  $\tau_j < \tau_{j-1} + e_{kij}$  was true, then we simulated  $e_{kij}$  using the  $j + 1$  interval parameters until either the inequality  $\tau_j \geq \tau_{j-1} + e_{kij}$  held for a particular  $j^*$  or the last time interval  $J$  was reached. We then defined  $T_{ki} = e_{kij^*} + \tau_{j^*-1}$ .

Censoring times  $\mathbf{C} = (\mathbf{C}_1^T, \dots, \mathbf{C}_k^T)^T$  where  $\mathbf{C}_k = (C_{k1}, \dots, C_{kn_k})^T$  were simulated from the uniform distribution  $C_{ki} \sim \text{Unif}(0, 5)$ , which assumes an end to follow-up after 5 years. The aforementioned event time simulation in combination with this censoring time simulation resulted in censoring rates that fell within 11% to 26% for all simulation conditions; average censoring rates across the 100 simulation replicates for the conditions mentioned ranged from 16% to 19%. The average

median follow-up time was approximately 0.60 years for the various conditions.

For individual  $i$  in group  $k$ , the vector of predictors for the fixed effects is given as  $\mathbf{x}_{ki} = (x_{ki,1}, \dots, x_{ki,p})^T$ , which does not include an intercept, and we define the random effects  $\mathbf{z}_{ki} = (1, \mathbf{x}_{ki})$ , where  $x_{ki,l} \sim N(0, 1)$  for  $l = 1, \dots, p$ , and each  $\mathbf{x}_i$  was standardized as described in Section 2.1. We include a random intercept in the random effects predictors  $\mathbf{z}_{ki}$  to allow for the baseline hazard to vary across groups.

We prepared the data to be fit with a piecewise constant hazard survival model by calculating eight time intervals—specified such that there were an approximately equal number of events within each time interval—and then creating the long-form dataset specified in Section 2.2 using the `survival::survSplit()` function from the **survival** R package [43, 44].

The results for these simulations are presented in Tables 1 and 2. Table 1 provides the average true and false positive percentages for both the fixed and random effects variable selection, the median time in hours to complete the variable selection procedure, the average of the mean absolute deviation between the fixed effects coefficient estimates and the true fixed effects coefficients across all simulation replicates, and the average of the Frobenius norm of the difference between the estimated random effect covariance matrix  $\hat{\mathbf{\Sigma}} = \hat{\mathbf{B}}\hat{\mathbf{B}}^T$  and the true covariance matrix  $\mathbf{\Sigma} = \mathbf{B}\mathbf{B}^T$  (the Frobenius norm was standardized by the number of random effects selected in the best model). The true positive percentages express the average percent of the true predictors selected in the best models across simulation replicates, and the false positive percentages express the average percent of false predictors selected in the best models. Table 2 gives the Growth Ratio estimation procedure results, including the average estimate of  $r$  and the proportion of times that the Growth Ratio estimate of  $r$  was underestimated, correct, or overestimated. All simulations were completed on a high performance computing cluster with CPU Intel processors between 2.3 and 2.5 GHz.

We see from Table 1 that the **phmmPen\_FA** method is able to accurately select both the fixed and random effects within the piecewise constant hazard mixed effects model across a variety of conditions. The true positive rates of the **phmmPen\_FA** method are generally above 90% for both fixed and random effects; the fixed effects true positives increase when the true predictor effects are larger, and the random effects true positives increase when the number of groups in the data increase. The false positive rates are less than 6.5% for fixed effects and less than 3.9% for the random effects across all conditions.

We can see from Table 2 that the Growth Ratio estimation procedure generally underestimates the number of latent factors  $r$  for the simulated data set-ups used in this section. We expect that this is a result of a combination of reasons, including relatively low numbers of groups  $K$  in the data and  $\mathbf{B}$  matrices that created  $\mathbf{\Sigma}$  matrices with relatively low eigenvalues. This is supported by the results that show an improvement in the accuracy of the estimation as the number of groups and the relative size of the  $\mathbf{B}$  matrix increases. Additionally, the Growth Ratio utilizes group-specific

**TABLE 1** | Variable selection results for the  $p = 100$  piecewise constant hazard mixed effects simulations, including true positive (TP) percentages for fixed and random effects, false positive (FP) percentages for fixed and random effects, the median time in hours for the algorithm to complete ( $T^{\text{med}}$ ), and the average of the mean absolute deviation (Abs. Dev. [Mean]) between the coefficient estimates and the true  $\beta$  values across all simulation replicates. Column **B** describes the general size of both the variances and eigenvalues of the resulting  $\Sigma = \mathbf{B}\mathbf{B}^T$  random effects covariance matrix. Column ‘ $r$  Est.’ refers to the method used to specify  $r$  in the algorithm: The Growth Ratio (GR) estimate or the true value of  $r$ . Column  $\|\mathbf{D}\|_F$  represents the average across simulation replicates of the Frobenius norm of the difference ( $\mathbf{D}$ ) between the estimated random effects covariance matrix  $\hat{\Sigma}$  and the true random effects covariance matrix  $\Sigma$ ; the Frobenius norm was standardized by the number of true random effects selected in the model.

$\beta$	$K$	<b>B</b>	$r$ est.	TP (%) Fixef	FP (%) Fixef	TP (%) Ranef	FP (%) Ranef	$T^{\text{med}}$	Abs. dev. (Mean)	$\ \mathbf{D}\ _F$
0.5	5	Small	True	91.80	2.39	93.00	0.65	3.81	0.23	0.31
			GR	90.80	2.46	91.60	1.40	2.34	0.22	0.42
		Moderate	True	91.00	4.85	94.40	1.31	6.96	0.34	0.63
			GR	91.00	4.18	92.40	2.13	3.78	0.32	0.70
	10	Small	True	94.60	2.18	98.60	0.99	4.81	0.17	0.27
			GR	94.00	3.28	95.80	1.63	2.66	0.17	0.31
		Moderate	True	90.00	5.25	94.60	3.57	6.08	0.24	0.54
			GR	86.20	6.42	93.40	3.83	3.02	0.23	0.61
1.0	5	Small	True	99.20	1.05	96.00	0.16	6.01	0.26	0.31
			GR	99.00	1.09	93.20	0.54	3.50	0.25	0.39
		Moderate	True	97.60	2.92	95.20	0.65	8.72	0.36	0.63
			GR	95.20	2.75	94.20	1.22	3.63	0.34	0.71
	10	Small	True	100.00	1.02	99.60	0.15	5.14	0.20	0.27
			GR	99.80	1.20	97.00	0.34	2.97	0.23	0.30
		Moderate	True	98.80	2.39	99.60	1.01	7.55	0.27	0.55
			GR	98.20	4.22	98.80	0.66	4.02	0.33	0.61

**TABLE 2** | Results of the Growth Ratio  $r$  estimation procedure for  $p = 100$  piecewise constant hazard mixed effects simulation results, including the average estimate of  $r$  across simulations and percent of times that the estimation procedure underestimated  $r$ , gave the true  $r$ , or overestimated  $r$ . Column **B** describes the general size of both the variances and eigenvalues of the resulting  $\Sigma = \mathbf{B}\mathbf{B}^T$  random effects covariance matrix.

$\beta$	$K$	<b>B</b>	Avg. $r$	$r$ underestimated (%)	$r$ correct (%)	$r$ overestimated (%)
0.5	5	Small	2.00	100	0	0
		Moderate	2.00	100	0	0
	10	Small	2.07	95	3	2
		Moderate	2.20	85	10	5
1.0	5	Small	2.00	100	0	0
		Moderate	2.00	100	0	0
	10	Small	2.13	88	11	1
		Moderate	2.19	83	15	2

penalized piecewise constant hazard coefficient estimates, and these estimates might be sensitive to the fact that for some simulated datasets, not all groups had a sufficient number of events within each time interval to get reasonable  $\tilde{\psi}$  estimates, possibly leading to less than stable pseudo random effect estimates. See a further discussion on this topic in Sections S4.1 and S3.5.

Even though the Growth Ratio procedure consistently underestimated  $r$ , this did not strongly impact the variable selection results nor the bias of the fixed effects estimates selected in the best models. When the algorithm used the Growth Ratio estimate of  $r$  instead of the true estimate of  $r$ , the true and false positive rates

remained very consistent, with only slight decreases in true positive rates for the fixed and random effects when the Growth Ratio procedure is used. The largest impact that underestimating  $r$  had on the bias of the fixed effects estimates was when  $K = 10$  and  $\beta = 1.0$ .

Our observation in these simulations that underestimating the number of latent factors  $r$  does not harm our method’s performance compared with using the true  $r$  can be explained by the eigenvalues of the random effects covariance matrix  $\Sigma$ . Our assumption that  $\Sigma$  can be represented by the low-rank (rank  $r$ ) matrix of  $\mathbf{B}\mathbf{B}^T$  (i.e., can be represented with  $r$  latent factors)



means that  $\Sigma$  has  $r$  non-zero eigenvalues. The sizes of these eigenvalues provide an indication of how important each latent factor is for representing  $\Sigma$ . Suppose we order these  $r$  eigenvalues from largest to smallest. If the first  $r - 1$  eigenvalues are relatively large and the  $r$ -th eigenvalue is relatively small, this means that most of the variation in  $\Sigma$  can be represented by the first  $r - 1$  latent factors (i.e., the first  $r - 1$  columns of the  $\mathbf{B}$  factor loadings matrix).

In our simulations, the moderate  $\mathbf{B}$  results in eigenvalues (3.38, 3.38, 2.25), and the small  $\mathbf{B}$  results in eigenvalues (1.5, 1.5, 1.0) (see details of the  $\mathbf{B}$  matrices in Section S1.1). Since the third eigenvalue is smaller than the first two, this indicates that the first two latent factors are the most important, and the third latent factor is of lesser importance. Therefore, if we underestimate  $r$  as 2 instead of 3 in our simulations, then we are perhaps not losing a large amount of information when we try and estimate the random effect covariance matrix  $\Sigma$ .

### 3.2 | Variable Selection in Survival Data With 500 Predictors

In order to further illustrate the scalability of our method, we applied our method to survival simulations with  $p = 500$  covariates. We simulated the event and censoring times from a piecewise constant hazard mixed effects model much like the procedure described in Section 3.1, except the total number of predictors used in the analyses was  $p = 500$  instead of  $p = 100$ . All simulations assumed the true number of latent factors  $r$  was 3 and the Growth Ratio method was used to estimate  $r$ . Just as in the  $p = 100$  survival simulations, we specified a full model for the algorithm such the random effect predictors equaled the fixed effect predictors (e.g.,  $q = p + 1$ ), and our aim was to select the set of true predictors and random effects. The variable selection results to these simulations are given in Table 3. The median times needed to complete these simulations took between 11.9 and 21.3 h.

When we compare specific data conditions (i.e., specific combinations of the size of the fixed and random effects and the number

of groups) between Tables 1 and 3, we see that increasing the number of total predictors input into the **phmmPen\_FA** procedure from  $p = 100$  to 500 tended to decrease the true positive rates for both the fixed and random effects and increase the total time needed to complete the variable selection procedure. Similar to the  $p = 100$  simulations, the Growth Ratio procedure continued to underestimate the number of latent factors in the underlying model.

### 3.3 | Supplemental Simulations

The Section S3 provides results from additional simulations. Section S3.1 compares our **phmmPen\_FA** method to a naive fixed-effects only variable selection approach (i.e., no random effects incorporated into the survival model), implemented using the **ncvreg** R package [25]. In terms of selecting the fixed effects predictors, the **phmmPen\_FA** method outperforms the naive fixed-effects only method of **ncvreg** when the true fixed effects coefficients are more moderate or when the true random effect coefficients are larger; **ncvreg** performs comparably to **phmmPen\_FA** in situations when the true fixed effects coefficients are larger and the random effects coefficients are smaller. As expected, the **ncvreg** package has the advantage of performing variable selection much faster since no random effects need to be selected.

Section S3.2 illustrates how the **phmmPen\_FA** method performs when the mixed effects survival data is simulated using the Weibull distribution instead of the piecewise constant hazard distribution. The results indicate that our method performs comparably for either data generation mechanism, and our method even performs slightly better on the Weibull-generated data according to some metrics.

Section S3.3 examines how the **phmmPen\_FA** method performs when the true underlying data has a different number of true fixed effects predictors and true random effects predictors. The results show that our method can also accurately select the fixed and random effects predictors in these scenarios.

**TABLE 3** | Variable selection results for the  $p = 500$  piecewise constant hazard mixed effects simulations, including true positive (TP) percentages for fixed and random effects, false positive (FP) percentages for fixed and random effects, the median time in hours for the algorithm to complete ( $T^{\text{med}}$ ), and the average of the mean absolute deviation (Abs. Dev. (Mean)) between the coefficient estimates and the true  $\beta$  values across all simulation replicates. Column **B** describes the general size of both the variances and eigenvalues of the resulting  $\Sigma = \mathbf{B}\mathbf{B}^T$  random effects covariance matrix. Column ' $r$  Est.' refers to the method used to specify  $r$  in the algorithm: The Growth Ratio (GR) estimate or the true value of  $r$ . Column  $\|\mathbf{D}\|_F$  represents the average across simulation replicates of the Frobenius norm of the difference ( $\mathbf{D}$ ) between the estimated random effects covariance matrix  $\hat{\Sigma}$  and the true random effects covariance matrix  $\Sigma$ ; the Frobenius norm was standardized by the number of true random effects selected in the model.

$\beta$	$K$	<b>B</b>	Avg. $r$	TP (%) Fixef	FP (%) Fixef	TP (%) Ranef	FP (%) Ranef	$T^{\text{med}}$	Abs. Dev. (Mean)	$\ \mathbf{D}\ _F$
0.5	5	Small	2.00	86.40	1.97	77.40	0.81	14.19	0.18	0.39
		Moderate	2.00	80.60	3.91	76.40	1.63	19.40	0.24	0.77
	10	Small	2.01	91.40	1.79	88.00	0.53	16.21	0.16	0.30
		Moderate	2.04	81.60	5.82	80.80	2.44	23.62	0.20	0.69
1.0	5	Small	2.00	99.40	0.36	91.40	0.02	18.20	0.30	0.30
		Moderate	2.00	93.60	0.79	85.00	0.12	19.77	0.39	0.72
	10	Small	2.06	100.00	0.56	95.40	0.04	16.78	0.30	0.27
		Moderate	2.03	97.20	1.31	92.80	0.18	24.31	0.40	0.64

Section S3.4 further explores how the **phmmPen\_FA** method performs when the number of latent factors  $r$  is purposefully underestimated. In general, decreasing  $r$  helps decrease the time needed to complete the **phmmPen\_FA** procedure. As the value of  $r$  used in the algorithm decreases and deviates further from the true value of  $r$ , the true positives of the fixed and random effects generally decreases; the false positives of the random effects generally increases; and the bias of the random effects covariance matrix generally increases. However, the differences between variable selection and bias performances across values of  $r$  continue to be fairly minor.

Section S3.5 investigates how the **phmmPen\_FA** method performs when different numbers of time intervals  $J$  are assumed (values of  $J$  from 5 to 10 were considered). In general, there was very little difference in the variable selection performance between different values of  $J$ . In the data scenarios considered in these simulations, smaller  $J$  values improved the accuracy of the Growth Ratio  $r$  estimate. Smaller  $J$  also decreased the median time needed to complete the procedure.

## 4 | Case Study: Pancreatic Ductal Adenocarcinoma

Patients diagnosed with Pancreatic Ductal Adenocarcinoma (PDAC) generally face a very poor prognosis, where the 5-year survival rate is 6% [45]. Consequently, it is of clinical interest to robustly identify gene signatures that are associated with overall survival to better predict patient prognosis in the clinic.

Selecting gene signatures for the prediction of clinical outcomes, including survival outcomes, can often be inconsistent across biomedical studies, where gene signatures identified in one study may have little or no overlap with ones identified in other studies [46]. Consequences of this lack of replicability in gene signature selection include variable accuracy in predicting clinical outcomes in new studies using these models [46, 47] and contradictory effect estimates relating genes to the outcome [48]. This lack of replicability across studies can come from small sample size [47] and differences in data pre-processing steps [49, 50], among other sources.

In order to improve replicability in the prediction of survival in PDAC, we combine PDAC gene expression data from seven different studies [51–57]. The studies used in these analyses are summarized in Table S1 (within Section S2.1). The seven combined studies resulted in a sample size of 879 subjects with 539 events. In order to account and adjust for between-study heterogeneity, we apply our new method **phmmPen\_FA** to fit a penalized piecewise constant hazard mixed effects model to our data to select predictors with study-replicable effects, where we assume that predictor effects may vary between studies.

Moffitt et al. [54] identified a 500-member gene list relevant to classifying two PDAC tumor subtypes they identified—basal and classical—which were prognostic of survival. Therefore, we decided to limit our initial interest to these 500 genes. Of these 500 genes, 420 of these genes were common among all of the datasets. We removed 20% of the genes with the lowest gene expression based on their average rank, leaving 336 genes.


We integrated gene expression data from multiple studies by first using the data integration rank transformation technique as specified by Rashid et al. [10], allowing us to sidestep complex questions regarding how to cross-normalize data. This integration technique creates top scoring pairs (TSPs). To illustrate the interpretation of TSPs, let  $g_{ki,A}$  and  $g_{ki,B}$  be the raw expression of genes  $A$  and  $B$  in subject  $i$  of group  $k$ . For each gene pair  $(g_{ki,A}, g_{ki,B})$ , the TSP is an indicator  $I(g_{ki,A} > g_{ki,B})$  which specifies which of the two genes has higher expression in the subject. We denote a TSP predictor as “GeneA\_GeneB”. In the dataset, we use 168 TSP predictors. The Section S2.1 provides additional details on the data processing and selection of the TSPs used in the analysis.

Due to the presence of several pairwise Spearman correlation values greater than 0.5 between the TSP covariates used within the analyses, we used the Elastic Net penalization procedure [24] to balance between ridge regression and the MCP penalty (regular MCP penalty [25] for the fixed effects, and grouped MCP penalty [26] for the rows of the  $\mathbf{B}$  matrix). We let  $\pi$  represent the balance between ridge regression and the MCP penalty, where  $\pi = 0$  represents ridge regression,  $\pi = 1$  represents the MCP penalty, and  $\pi \in (0, 1)$  represents a combination of the two.

We used the **phmmPen\_FA** procedure to fit this PDAC survival data with a penalized piecewise constant hazard mixed effect survival model. Like in the simulations presented in Section 3, we assumed  $J = 8$  time intervals, with the interval boundaries selected such that an approximately equal number of events within each time interval. We considered values of  $\pi = \{0.7, 0.8, 0.9, 1.0\}$ , where the same value of  $\pi$  was used for both the fixed and random effects penalization, and  $r$  evaluated using the Growth Ratio procedure (evaluated as 2 for all cases) and  $r$  manually set to a larger value of 3 since our simulations indicated that the GR estimate of  $r$  may be underestimated. The sequence of  $\lambda$  penalties used in the variable selection procedure is described Section S1.3. The best model within each  $\pi$  and  $r$  combination was selected using the BIC-ICQ model selection criteria.

To evaluate the performance of **phmmPen\_FA** under the various  $\pi$  and  $r$  combinations described above, we utilized cross validation to estimate the C-index for each set-up. We randomly selected 80% of the observations from each study to be the training dataset (with random selection stratified by events and censored observations), and the remaining 20% of the observations were set as the test dataset. The C-index was calculated using the **intsurv::cIndex** function [58], and the C-index risk score was calculated as the ‘best’ model’s fixed effects coefficients applied to the training dataset TSP predictors; the absence of any random effect coefficients in this risk score calculation was done in order to replicate typical real-world mixed effects scenarios, where the groups within the data used to train the models will not often equal the groups in future data to which the model will be applied.

The combination of  $\pi = 0.9$  and  $r = 3$  produced the largest C-index value of 0.6511. This combination selected 15 of the 168 total TSPs to have non-zero fixed effects (see Figure 1) and were therefore considered important for the prediction of survival in PDAC subjects. One TSP, CBLC\_SMURF1, was selected to have a non-zero random effect (random effect variance estimated as 0.035). The time to complete the algorithm was 2.1 h. The overall



**FIGURE 1** | Log hazard ratios for the TSP covariates selected during the case study analysis.

range of the C-index values across the various  $\pi$  and  $r$  conditions was fairly small (smallest C-index value was 0.6445 for  $\pi = 0.7$  and  $r = 2$  from the GR estimation procedure). This small range for the C-index is likely a consequence of the high proportion of censored observations within this dataset [59].

Further details about the sensitivity analyses for this case study can be found in the Section S2.2.

## 5 | Discussion

We have shown through simulations and a case study of pancreatic ductal adenocarcinoma patients that we can extend the method to perform variable selection in high dimensional mixed effects models to survival data. We accomplish this by approximating proportional hazards mixed effects models using a piecewise constant hazard mixed effects model and then applying the MCECM algorithm to simultaneously select for fixed and random effects. We incorporate the factor model decomposition of the random effects proposed in Heiling et al. [13] in order to scale this method to larger dimensions, for example, hundreds of predictors.

The simulations presented in Section 3 show that the **phmm-Pen\_FA** method can accurately select both fixed and random effects even for small or moderate effect sizes, which reflects hazard values and variations in typical survival data. By using the factor model decomposition of the random effects, this model selection procedure can be accomplished within reasonable time frames even when we consider hundreds of predictors as input for the variable selection procedure.

Our method is limited by the need to provide an estimate for the number of latent factors that model the random effects. The simulation results showed that the Growth Ratio procedure tended to underestimate this value for the simulation conditions that we considered. However, even when the number of latent factors

was estimated incorrectly by the Growth Ratio procedure, this mis-specification had very little impact on the general variable selection performance or the fixed effects coefficient estimates. Therefore, our method is not sensitive to the estimation of the number of latent factors.

Similar to other penalization approaches to variable selection such as the **glmnet** [24] and **ncvreg** [25] R packages, another limitation of our method is that it focuses on selecting relevant predictors but does not provide inference or measurements of uncertainty for the selected predictors. In order to get inference or uncertainty measures, one would have to take the selected model from the **phmmPen\_FA** procedure and use other software that could fit mixed effects models to provide this information, such as the **coxme** R package [60] that fits survival mixed effects models.

In our simulations in Section 3 and our case study analysis in Section 4, we have not compared the performance of our **phmm-Pen\_FA** method with other proportional hazards mixed effects (PHMM) models because we were not aware of other PHMM methods that could select both fixed and random effects and be applied to survival data with hundreds of predictors. However, we do provide a comparison between our method and a naive method that performs variable selection on fixed effects only using the **ncvreg** R package, see Section S3.1 for these results.

### Conflicts of Interest

The authors declare no conflicts of interest.

### Data Availability Statement

The data that support the findings of this study are openly available in the GitHub repository [https://github.com/hheiling/paper\\_phmmPen\\_FA](https://github.com/hheiling/paper_phmmPen_FA).

### References

1. F. Vaida and R. Xu, “Proportional Hazards Model With Random Effects,” *Statistics in Medicine* 19, no. 24 (2000): 3309–3324.

2. S. Ripatti and J. Palmgren, "Estimation of Multivariate Frailty Models Using Penalized Partial Likelihood," *Biometrics* 56, no. 4 (2000): 1016–1022.
3. R. Tibshirani, "The Lasso Method for Variable Selection in the Cox Model," *Statistics in Medicine* 16, no. 4 (1997): 385–395.
4. J. Bradic, J. Fan, and J. Jiang, "Regularization for Cox's Proportional Hazards Model With NP-Dimensionality," *Annals of Statistics* 39, no. 6 (2011): 3092.
5. N. Simon, J. Friedman, T. Hastie, and R. Tibshirani, "Regularization Paths for Cox's Proportional Hazards Model via Coordinate Descent," *Journal of Statistical Software* 39, no. 5 (2011): 1–13, <https://doi.org/10.18637/jss.v039.i05>.
6. J. Fan and R. Li, "Variable Selection for Cox's Proportional Hazards Model and Frailty Model," *Annals of Statistics* 30, no. 1 (2002): 74–99.
7. R. Xu, F. Vaida, and D. P. Harrington, "Using Profile Likelihood for Semiparametric Model Selection With Application to Proportional Hazards Mixed Models," *Statistica Sinica* 19, no. 2 (2009): 819–842.
8. M. Donohue, R. Overholser, R. Xu, and F. Vaida, "Conditional Akaike Information Under Generalized Linear and Proportional Hazards Mixed Models," *Biometrika* 98, no. 3 (2011): 685–700.
9. K. E. Lee, Y. Kim, and R. Xu, "Bayesian Variable Selection Under the Proportional Hazards Mixed-Effects Model," *Computational Statistics and Data Analysis* 75 (2014): 53–65.
10. N. U. Rashid, Q. Li, J. J. Yeh, and J. G. Ibrahim, "Modeling Between-Study Heterogeneity for Improved Replicability in Gene Signature Selection and Clinical Prediction," *Journal of the American Statistical Association* 115, no. 531 (2020): 1125–1138.
11. H. Heiling, N. Rashid, Q. Li, and J. Ibrahim, "GlmmPen: High Dimensional Penalized Generalized Linear Mixed Models (pGLMM)," 2023 R package version 1.5.3.4, <https://cran.r-project.org/package=glmmPen>.
12. H. M. Heiling, N. U. Rashid, Q. Li, and J. G. Ibrahim, "GlmmPen: High Dimensional Penalized Generalized Linear Mixed Models," *R Journal* 15, no. 4 (2023): 106–128, <https://doi.org/10.32614/rj-2023-086>.
13. H. M. Heiling, N. U. Rashid, Q. Li, X. L. Peng, J. J. Yeh, and J. G. Ibrahim, "Efficient Computation of High-Dimensional Penalized Generalized Linear Mixed Models by Latent Factor Modeling of the Random Effects," *Biometrics* 80, no. 1 (2024): ujae016, <https://doi.org/10.1093/biomtc/ujae016>.
14. P. D. Allison, *Survival Analysis Using SAS: A Practical Guide* (Cary, NC: SAS Institute, 2010).
15. T. R. Holford, "The Analysis of Rates and of Survivorship Using Log-Linear Models," *Biometrics* 36, no. 2 (1980): 299–305.
16. N. Laird and D. Olivier, "Covariance Analysis of Censored Survival Data Using Log-Linear Analysis Techniques," *Journal of the American Statistical Association* 76, no. 374 (1981): 231–240.
17. M. Friedman, "Piecewise Exponential Models for Survival Data With Covariates," *Annals of Statistics* 10, no. 1 (1982): 101–113.
18. J. G. Ibrahim, M. H. Chen, and D. Sinha, *Bayesian Survival Analysis*. 2 (New York, NY: Springer, 2001).
19. P. C. Austin, "A Tutorial on Multilevel Survival Analysis: Methods, Models and Applications," *International Statistical Review* 85, no. 2 (2017): 185–203.
20. R. J. Gray, "A Bayesian Analysis of Institutional Effects in a Multicenter Cancer Clinical Trial," *Biometrics* 50, no. 1 (1994): 244–253.
21. J. G. Ibrahim, M. H. Chen, and S. N. MacEachern, "Bayesian Variable Selection for Proportional Hazards Models," *Canadian Journal of Statistics* 27, no. 4 (1999): 701–717.
22. Z. Chen and D. B. Dunson, "Random Effects Selection in Linear Mixed Models," *Biometrics* 59, no. 4 (2003): 762–769.
23. J. G. Ibrahim, H. Zhu, R. I. Garcia, and R. Guo, "Fixed and Random Effects Selection in Mixed Effects Models," *Biometrics* 67, no. 2 (2011): 495–503.
24. J. Friedman, T. Hastie, and R. Tibshirani, "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software* 33, no. 1 (2010): 1–22.
25. P. Breheny and J. Huang, "Coordinate Descent Algorithms for Non-convex Penalized Regression, With Applications to Biological Feature Selection," *Annals of Applied Statistics* 5, no. 1 (2011): 232–253.
26. P. Breheny and J. Huang, "Group Descent Algorithms for Nonconvex Penalized Linear and Logistic Regression Models With Grouped Predictors," *Statistics and Computing* 25, no. 2 (2015): 173–187.
27. R. I. Garcia, J. G. Ibrahim, and H. Zhu, "Variable Selection for Regression Models With Missing Data," *Statistica Sinica* 20, no. 1 (2010): 149–165.
28. B. Carpenter, A. Gelman, M. D. Hoffman, et al., "Stan: A Probabilistic Programming Language," *Journal of Statistical Software* 28 (2017): 1–32, <https://doi.org/10.18637/jss.v076.i01>.
29. M. D. Hoffman and A. Gelman, "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research* 15, no. 1 (2014): 1593–1623.
30. Stan Development Team, "RStan: The R interface to Stan," 2020 R package version 2.21.2, <https://CRAN.R-project.org/package=rstan>.
31. G. O. Roberts and J. S. Rosenthal, "Examples of Adaptive MCMC," *Journal of Computational and Graphical Statistics* 18, no. 2 (2009): 349–367.
32. G. H. Givens and J. A. Hoeting, *Computational Statistics*, 2nd ed. (Hoboken, NJ: John Wiley & Sons, 2012), 703.
33. N. Parikh and S. Boyd, "Proximal Algorithms," *Foundations and Trends® in Optimization* 1, no. 3 (2014): 127–239.
34. S. C. Ahn and A. R. Horenstein, "Eigenvalue Ratio Test for the Number of Factors," *Econometrica* 81, no. 3 (2013): 1203–1227.
35. SAS Institute Inc, *SAS/STAT Software, Version 9.4 15.1 User's Guide* (Cary, NC: SAS Institute, Inc., 2018).
36. I. Castillo and v. Pas, "Multiscale Bayesian Survival Analysis," *Annals of Statistics* 49, no. 6 (2021): 3559–3582.
37. Y. Li, M. H. Gail, D. L. Preston, B. I. Graubard, and J. H. Lubin, "Piecewise Exponential Survival Times and Analysis of Case-Cohort Data," *Statistics in Medicine* 31, no. 13 (2012): 1361–1368.
38. Z. Qiou, N. Ravishanker, and D. K. Dey, "Multivariate Survival Analysis With Positive Stable Frailties," *Biometrics* 55, no. 2 (1999): 637–644.
39. D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, "Bayesian Measures of Model Complexity and Fit," *Journal of the Royal Statistical Society, Series B: Statistical Methodology* 64, no. 4 (2002): 583–639.
40. V. G. Cancho, J. Rodrigues, and M. Castro, "A Flexible Model for Survival Data With a Cure Rate: A Bayesian Approach," *Journal of Applied Statistics* 38, no. 1 (2011): 57–70.
41. J. G. Ibrahim, M. H. Chen, and H. Chu, "Bayesian Methods in Clinical Trials: A Bayesian Analysis of ECOG Trials E1684 and E1690," *BMC Medical Research Methodology* 12, no. 1 (2012): 1–12.
42. J. G. Ibrahim, M. H. Chen, and D. Sinha, "Criterion-Based Methods for Bayesian Model Assessment," *Statistica Sinica* 11, no. 2 (2001): 419–443.
43. T. M. Therneau, "A Package for Survival Analysis in R," 2021 R package version 3.2-11, <https://CRAN.R-project.org/package=survival>.
44. T. M. Therneau and P. M. Grambsch, *Modeling Survival Data: Extending the Cox Model* (New York: Springer, 2000).

45. A. A. Khorana, P. B. Mangu, J. Berlin, et al., “Potentially Curable Pancreatic Cancer: American Society of Clinical Oncology Clinical Practice Guideline,” *Journal of Clinical Oncology* 34, no. 21 (2016): 2541–2556.
46. L. Waldron, B. Haibe-Kains, A. C. Culhane, et al., “Comparative Meta-Analysis of Prognostic Gene Signatures for Late-Stage Ovarian Cancer,” *JNCI Journal of the National Cancer Institute* 106, no. 5 (2014): 1–12.
47. C. Sotiriou and M. J. Piccart, “Taking Gene-Expression Profiling to the Clinic: When Will Molecular Signatures Become Relevant to Patient Care?,” *Nature Reviews Cancer* 7, no. 7 (2007): 545–553.
48. E. M. Swisher, T. Taniguchi, and B. Y. Karlan, “Molecular Scores to Predict Ovarian Cancer Outcomes: A Worthy Goal, But Not Ready for Prime Time,” 2012.
49. L. Lusa, L. M. McShane, J. F. Reid, et al., “Challenges in Projecting Clustering Results Across Gene Expression–Profiling Datasets,” *JNCI Journal of the National Cancer Institute* 99, no. 22 (2007): 1715–1723.
50. E. R. Paquet and M. T. Hallett, “Absolute Assignment of Breast Cancer Intrinsic Molecular Subtype,” *Journal of the National Cancer Institute* 107, no. 1 (2015): dju357, <https://doi.org/10.1093/jnci/dju357>.
51. A. J. Aguirre, J. A. Nowak, N. D. Camarda, et al., “Real-Time Genomic Characterization of Advanced Pancreatic Cancer to Enable Precision Medicine,” *Cancer Discovery* 8, no. 9 (2018): 1096–1111.
52. L. Cao, C. Huang, D. C. Zhou, et al., “Proteogenomic Characterization of Pancreatic Ductal Adenocarcinoma,” *Cell* 184, no. 19 (2021): 5031–5052.
53. F. Dijk, V. L. Veenstra, E. C. Soer, et al., “Unsupervised Class Discovery in Pancreatic Ductal Adenocarcinoma Reveals Cell-Intrinsic Mesenchymal Features and High Concordance Between Existing Classification Systems,” *Scientific Reports* 10, no. 1 (2020): 337, <https://doi.org/10.1038/s41598-019-56826-9>.
54. R. A. Moffitt, R. Marayati, E. L. Flate, et al., “Virtual Microdissection Identifies Distinct Tumor-and Stroma-Specific Subtypes of Pancreatic Ductal Adenocarcinoma,” *Nature Genetics* 47, no. 10 (2015): 1168–1178.
55. P. Bailey, D. K. Chang, K. Nones, et al., “Genomic Analyses Identify Molecular Subtypes of Pancreatic Cancer,” *Nature* 531, no. 7592 (2016): 47–52.
56. F. Puleo, R. Nicolle, Y. Blum, et al., “Stratification of Pancreatic Ductal Adenocarcinomas Based on Tumor and Microenvironment Features,” *Gastroenterology* 155, no. 6 (2018): 1999–2013.
57. B. J. Raphael, R. H. Hruban, A. J. Aguirre, et al., “Integrated Genomic Characterization of Pancreatic Ductal Adenocarcinoma,” *Cancer Cell* 32, no. 2 (2017): 185–203.
58. W. Wang, K. Chen, and J. Yan, “*Intsurv: Integrative Survival Models*,” 2021 R package version 0.2.2, <https://CRAN.R-project.org/package=intsurv>.
59. A. Alabdallah, M. Ohlsson, S. Pashami, and T. Rögnvaldsson, “The Concordance Index Decomposition: A Measure for a Deeper Understanding of Survival Prediction Models,” *Artificial Intelligence in Medicine* 148 (2024): 102781, <https://doi.org/10.1016/j.artmed.2024.102781>.
60. T. M. Therneau, “Coxme: Mixed Effects Cox Models,” 2022 R package version 2.2-18.1.

## Supporting Information

Additional supporting information can be found online in the Supporting Information section.