

C-mix: A high-dimensional mixture model for censored durations, with applications to genetic data

Simon Bussy,^{1,2}  Agathe Guilloux,³ Stéphane Gaïffas^{4,5} and Anne-Sophie Jannot^{6,7}

Statistical Methods in Medical Research
0(0) 1–17

© The Author(s) 2018

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0962280218766389

journals.sagepub.com/home/smm



Abstract

We introduce a supervised learning mixture model for censored durations (C-mix) to simultaneously detect subgroups of patients with different prognosis and order them based on their risk. Our method is applicable in a high-dimensional setting, i.e. with a large number of biomedical covariates. Indeed, we penalize the negative log-likelihood by the Elastic-Net, which leads to a sparse parameterization of the model and automatically pinpoints the relevant covariates for the survival prediction. Inference is achieved using an efficient Quasi-Newton Expectation Maximization algorithm, for which we provide convergence properties. The statistical performance of the method is examined on an extensive Monte Carlo simulation study and finally illustrated on three publicly available genetic cancer datasets with high-dimensional covariates. We show that our approach outperforms the state-of-the-art survival models in this context, namely both the CURE and Cox proportional hazards models penalized by the Elastic-Net, in terms of C-index, $AUC(t)$ and survival prediction. Thus, we propose a powerful tool for personalized medicine in cancerology.

Keywords

Coxs proportional hazards model, CURE model, Elastic-Net regularization, high-dimensional estimation, mixture duration model, survival analysis

1 Introduction

Predicting subgroups of patients with different prognosis is a key challenge for personalized medicine, see for instance Alizadeh et al.¹ and Rosenwald et al.² where subgroups of patients with different survival rates are identified based on gene expression data. A substantial number of techniques can be found in the literature to predict the subgroup of a given patient in a classification setting, namely when subgroups are known in advance.^{3,4,5} We consider in the present paper the much more difficult case where subgroups are unknown.

In this situation, a first widespread approach consists in first using unsupervised learning techniques applied on the covariates – for instance, on the gene expression data^{6,7,8} – to define subsets of patients and then estimating the risks in each of them. The problem of such techniques is that there is no guarantee that the identified subgroups will have different risks. Another approach to subgroups identification is conversely based exclusively on the survival times: patients are then assigned to a low-risk or a high-risk group based on whether they were still alive.^{9,10} The problem here is that the resulting subgroups may not be biologically meaningful since the method do not use the covariates, and prognosis prediction based on covariates is not possible.

¹Theoretical and Applied Statistics Laboratory, Pierre and Marie Curie University, Paris, France

²Laboratoire de Probabilités Statistique et Modélisation (LPSM), UMR 8001, Sorbonne University, Paris, France

³LAMME, Univ Evry, CNRS, Université Paris-Saclay, Evry, France

⁴Center for Applied Mathematics, Ecole Polytechnique, Palaiseau, France

⁵LPMA, UMR CNRS 7599, Paris Diderot University, Paris, France

⁶Assistance Publique-Hôpitaux de Paris, Biomedical Informatics and Public Health Department, European Georges Pompidou Hospital, Paris, France

⁷INSERM UMRS 1138, Eq22, Centre de Recherche des Cordeliers, Université Paris Descartes, Paris, France

Corresponding author:

Simon Bussy, Laboratoire de Probabilités Statistique et Modélisation (LPSM), UMR 8001, Sorbonne University, Paris, France.

Email: simon.bussy@gmail.com

The method we propose uses both the survival information of the patients and its covariates in a supervised learning way. Moreover, it relies on the idea that exploiting the subgroups structure of the data, namely the fact that a portion of the population have a higher risk of early death, could improve the survival prediction of future patients (unseen during the learning phase).

We propose to consider a mixture of event times distributions in which the probabilities of belonging to each subgroups are driven by the covariates (e.g. gene expression data, patients characteristics, therapeutic strategy or omics covariates). Our C-mix model is hence part of the class of model-based clustering algorithms, as introduced in Banfield and Raftery.¹¹

More precisely, to model the heterogeneity within the patient population, we introduce a latent variable $Z \in \{0, \dots, K-1\}$ and our focus is on the conditional distribution of Z given the values of the covariates $X=x$. Now, conditionally on the latent variable Z , the distribution of duration time T is different, leading to a mixture in the event times distribution.

For a patient with covariates x , the conditional probabilities $\pi_k(x) = \mathbb{P}[Z = k|X = x]$ of belonging to the k -th risk group can be seen as scores that can help in decision-making for physicians. As a byproduct, it can also shed light on the effect of the covariates (which combination of biomedical markers is relevant to a given event of interest).

Our methodology differs from the standard survival analysis approaches in various ways, which we describe in this paragraph. First, the Cox proportional hazards (PH) model¹² (by far the most widely used in such a setting) is a regression model that describes the relation between intensity of events and covariates x via

$$\lambda(t) = \lambda_0(t)\exp(x^\top \beta^{\text{cox}}) \quad (1)$$

where λ_0 is a baseline intensity, and β^{cox} is a vector quantifying the multiplicative impact on the hazard ratio of each covariate. As in our model, high-dimensional covariates can be handled, through e.g. penalization, see Simon et al.¹³ But it does not permit the stratification of the population in groups of homogeneous risks, hence does not deliver a simple tool for clinical practice. Moreover, we show in the numerical sections that the C-mix model can be trained very efficiently in high dimension, and outperforms the standard Cox PH model by far in the analysed datasets.

Other models consider mixtures of event times distributions. In the CURE model (see Farewell¹⁴ and Kuk and Chen¹⁵), one fraction of the population is considered as cured (hence not subject to any risk). This can be very limiting, as for a large number of applications (e.g. rehospitalization for patients with chronic diseases or relapse for patients with metastatic cancer), all patients are at risk. We consider, in our model, that there is always an event risk, no matter how small. Other mixture models have been considered in survival analysis: see Kuo and Peng¹⁶ for a general study about mixture model for survival data or De Angelis et al.¹⁷ in a cancer survival analysis setting, to name but a few. Unlike our algorithm, none of these algorithms considers the high-dimensional setting.

A precise description of the model is given in Section 2. Section 3 focuses on the regularized version of the model with an Elastic-Net penalization to exploit dimension reduction and prevent overfitting. Inference is presented under this framework, as well as the convergence properties of the developed algorithm. Section 4 highlights the simulation procedure used to evaluate the performances and compares it with state-of-the-art models. In Section 5, we apply our method to genetic datasets. Finally, we discuss the obtained results in Section 6.

2 A censored mixture model

Let us present the survival analysis framework. We assume that the conditional density of the duration T given $X=x$ is a mixture

$$f(t|X=x) = \sum_{k=0}^{K-1} \pi_k(x) f_k(t; \alpha_k)$$

of $K \geq 1$ densities f_k , for $t \geq 0$ and $\alpha_k \in \mathbb{R}^{d_k}$ some parameters to estimate. The weights combining these distributions depend on the patient biomedical covariates x and are such that

$$\sum_{k=0}^{K-1} \pi_k(x) = 1 \quad (2)$$

This is equivalent to saying that conditionally on a latent variable $Z = k \in \{0, \dots, K-1\}$, the density of T at time $t \geq 0$ is $f_k(t; \alpha_k)$, and we have

$$\mathbb{P}[Z = k|X = x] = \pi_k(x) = \pi_{\beta_k}(x)$$

where $\beta_k = (\beta_{k,1}, \dots, \beta_{k,d}) \in \mathbb{R}^d$ denotes a vector of coefficients that quantifies the impact of each biomedical covariates on the probability that a patient belongs to the k -th group. Consider a logistic link function for these weights given by

$$\pi_{\beta_k}(x) = \frac{e^{x^\top \beta_k}}{\sum_{k=0}^{K-1} e^{x^\top \beta_k}} \quad (3)$$

The hidden status Z has therefore a multinomial distribution $\mathcal{M}(\pi_{\beta_0}(x), \dots, \pi_{\beta_{K-1}}(x))$. The intercept term is here omitted without loss of generality.

In practice, information loss occurs of right censoring type. This is taken into account in our model by introducing the following: a time $C \geq 0$ when the individual ‘leaves’ the target cohort, a right-censored duration Y and a censoring indicator Δ , defined by

$$Y = \min(T, C) \quad \text{and} \quad \Delta = \mathbb{1}_{\{T \leq C\}}$$

where $\min(a, b)$ denotes the minimum between two numbers a and b , and $\mathbb{1}$ denotes the indicator function.

In order to write a likelihood and draw inference, we make the two following hypothesis.

Hypothesis 1: T and C are conditionally independent given Z and X .

Hypothesis 2: C is independent of Z .

Hypothesis 1 is classical in survival analysis,¹⁸ while Hypothesis 2 is classical in survival mixture models.^{16,17} Under this hypothesis, denoting g the density of the censoring C , F the cumulative distribution function corresponding to a given density f , $\bar{F} = 1 - F$ and $F(y^-) = \lim_{u \rightarrow y^-} F(u)$, we have

$$\begin{aligned} \mathbb{P}[Y \leq y, \Delta = 1] &= \mathbb{P}[T \leq y, T \leq C] = \int_0^y f(u) \bar{G}(u) du \quad \text{and} \\ \mathbb{P}[Y \leq y, \Delta = 0] &= \mathbb{P}[C \leq y, C < T] = \int_0^y g(u) \bar{F}(u) du \end{aligned}$$

Then, denoting $\theta = (\alpha_0, \dots, \alpha_{K-1}, \beta_0, \dots, \beta_{K-1})^\top$ the parameters to infer and considering an independent and identically distributed (i.i.d.) cohort of n patients $(x_1, y_1, \delta_1), \dots, (x_n, y_n, \delta_n) \in \mathbb{R}^d \times \mathbb{R}_+ \times \{0, 1\}$, the log-likelihood of the C-mix model can be written

$$\begin{aligned} \ell_n(\theta) = \ell_n(\theta; \mathbf{y}, \boldsymbol{\delta}) &= n^{-1} \sum_{i=1}^n \left\{ \delta_i \log \left[\bar{G}(y_i^-) \sum_{k=0}^{K-1} \pi_{\beta_k}(x_i) f_k(y_i; \alpha_k) \right] \right. \\ &\quad \left. + (1 - \delta_i) \log \left[g(y_i) \sum_{k=0}^{K-1} \pi_{\beta_k}(x_i) \bar{F}_k(y_i^-; \alpha_k) \right] \right\} \end{aligned}$$

where we use the notations $\mathbf{y} = (y_1, \dots, y_n)^\top$ and $\boldsymbol{\delta} = (\delta_1, \dots, \delta_n)^\top$. Note that from now on, all computations are done conditionally on the covariates $(x_i)_{i=1, \dots, n}$. An important fact is that we *do not need to know or parameterize* \bar{G} nor g , namely the distribution of the censoring, for inference in this model (since all \bar{G} and g terms vanish in equation (6)).

3 Inference of C-mix

In this section, we describe the procedure for estimating the parameters of the C-mix model. We begin by presenting the Quasi-Newton Expectation Maximization (QNEM) algorithm we use for inference. We then focus our study on the convergence properties of the algorithm.

3.1 QNEM algorithm

In order to avoid overfitting and to improve the prediction power of our model, we use Elastic-Net regularization¹⁹ by minimizing the penalized objective

$$\ell_n^{\text{pen}}(\theta) = -\ell_n(\theta) + \sum_{k=0}^{K-1} \gamma_k \left((1 - \eta) \|\beta_k\|_1 + \frac{\eta}{2} \|\beta_k\|_2^2 \right) \quad (4)$$

where we add a linear combination of the lasso (ℓ_1) and ridge (squared ℓ_2) penalties for a fixed $\eta \in [0, 1]$, tuning parameter γ_k , and where we denote $\|\beta_k\|_p = (\sum_{i=1}^d |\beta_{k,i}|^p)^{1/p}$ the ℓ_p -norm of β_k . One advantage of this regularization method is its ability to perform model selection (the lasso part) and pinpoint the most important covariates relatively to the prediction objective. On the other hand, the ridge part allows to handle potential correlation between covariates.¹⁹ Note that in practice, the intercept is not regularized.

In order to derive an algorithm for this objective, we introduce a so-called QNEM, being a combination between an EM algorithm²⁰ and a L-BFGS-B algorithm.²¹ For the EM part, we need to compute the negative completed log-likelihood (here scaled by n^{-1}), namely the negative joint distribution of \mathbf{y} , $\boldsymbol{\delta}$ and $\mathbf{z} = (z_1, \dots, z_n)^\top$. It can be written

$$\begin{aligned} \ell_n^{\text{comp}}(\theta) &= \ell_n^{\text{comp}}(\theta; \mathbf{y}, \boldsymbol{\delta}, \mathbf{z}) \\ &= -n^{-1} \sum_{i=1}^n \left\{ \delta_i \left[\sum_{k=0}^{K-1} \mathbb{1}_{\{z_i=k\}} (\log \pi_{\beta_k}(x_i) + \log f_k(y_i; \alpha_k)) + \log \bar{G}(y_i^-) \right] \right. \\ &\quad \left. + (1 - \delta_i) \left[\sum_{k=0}^{K-1} \mathbb{1}_{\{z_i=k\}} (\log \pi_{\beta_k}(x_i) + \log \bar{F}_k(y_i^-; \alpha_k)) + \log g(y_i) \right] \right\} \end{aligned} \quad (5)$$

Suppose that we are at step $l+1$ of the algorithm, with current iterate denoted $\theta^{(l)} = (\alpha_0^{(l)}, \dots, \alpha_{K-1}^{(l)}, \beta_0^{(l)}, \dots, \beta_{K-1}^{(l)})^\top$. For the E-step, we need to compute the expected log-likelihood given by

$$Q_n(\theta, \theta^{(l)}) = \mathbb{E}_{\theta^{(l)}}[\ell_n^{\text{comp}}(\theta) | \mathbf{y}, \boldsymbol{\delta}]$$

We note that

$$q_{i,k}^{(l)} = \mathbb{E}_{\theta^{(l)}}[\mathbb{1}_{\{z_i=k\}} | y_i, \delta_i] = \mathbb{P}_{\theta^{(l)}}[z_i = k | y_i, \delta_i] = \frac{\Lambda_{k,i}^{(l)}}{\sum_{r=0}^{K-1} \Lambda_{r,i}^{(l)}} \quad (6)$$

with

$$\Lambda_{k,i}^{(l)} = [f_k(y_i; \alpha_k^{(l)}) \bar{G}(y_i^-)]^{\delta_i} [g(y_i) \bar{F}_k(y_i^-; \alpha_k^{(l)})]^{1-\delta_i} \pi_{\beta_k^{(l)}}(x_i) \quad (7)$$

so that $Q_n(\theta, \theta^{(l)})$ is obtained from equation (5) by replacing the two $\mathbb{1}_{\{z_i=k\}}$ occurrences with $q_{i,k}^{(l)}$. Depending on the chosen distributions f_k , the M-step can either be explicit for the updates of α_k (see Section 3.3 below for the geometric distributions case), or obtained using a minimization algorithm otherwise.

Let us focus now on the update of β_k in the M-step of the algorithm. By denoting

$$R_{n,k}^{(l)}(\beta_k) = -n^{-1} \sum_{i=1}^n q_{i,k}^{(l)} \log \pi_{\beta_k}(x_i)$$

the quantities involved in Q_n that depend on β_k , the update for β_k therefore requires to minimize

$$R_{n,k}^{(l)}(\beta_k) + \gamma_k \left((1 - \eta) \|\beta_k\|_1 + \frac{\eta}{2} \|\beta_k\|_2^2 \right) \quad (8)$$

The minimization Problem (8) is a convex problem. It looks like the logistic regression objective, where labels are not fixed but softly encoded by the expectation step (computation of $q_{i,k}^{(l)}$ above, see equation (6)).

We minimize equation (8) using the well-known L-BFGS-B algorithm.²¹ This algorithm belongs to the class of quasi-Newton optimization routines, which solve 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 penalization, which is not differentiable, we use the trick borrowed from Andrew and Gao:²² 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 (8) as the following differentiable problem with box constraints

$$\begin{aligned} & \text{minimize} && R_{n,k}^{(l)}(\beta_k^+ - \beta_k^-) + \gamma_k(1 - \eta) \sum_{j=1}^d (\beta_{k,j}^+ + \beta_{k,j}^-) + \gamma_k \frac{\eta}{2} \|\beta_k^+ - \beta_k^-\|_2^2 \\ & \text{subject to} && \beta_{k,j}^+ \geq 0 \quad \text{and} \quad \beta_{k,j}^- \geq 0 \quad \text{for all } j \in \{1, \dots, d\} \end{aligned} \quad (9)$$

where $\beta_k^\pm = (\beta_{k,1}^\pm, \dots, \beta_{k,d}^\pm)^\top$. The L-BFGS-B solver requires the exact value of the gradient, which is easily given by

$$\frac{\partial R_{n,k}^{(l)}(\beta_k)}{\partial \beta_k} = -n^{-1} \sum_{i=1}^n q_{i,k}^{(l)} (1 - \pi_{\beta_k}(x_i)) x_i \quad (10)$$

In Algorithm 1, we describe the main steps of the QNEM algorithm to minimize the function given in equation (4).

Algorithm 1: QNEM Algorithm for inference of the C-mix model

Require. Training data $(x_i, y_i, \delta_i)_{i \in \{1, \dots, n\}}$; starting parameters $(\alpha_k^{(0)}, \beta_k^{(0)})_{k \in \{0, \dots, K-1\}}$; tuning parameters $\gamma_k \geq 0$.

- 1: **for** $l = 0, \dots$, until convergence **do**
 - 2: Compute $(q_{i,k}^{(l)})_{k \in \{0, \dots, K-1\}}$ using equation (6).
 - 3: Compute $(\alpha_k^{(l+1)})_{k \in \{0, \dots, K-1\}}$.
 - 4: Compute $(\beta_k^{(l+1)})_{k \in \{0, \dots, K-1\}}$ by solving equation (9) with the L-BFGS-B algorithm.
 - 5: **end for**
 - 6: **return** Last parameters $(\alpha_k^{(l)}, \beta_k^{(l)})_{k \in \{0, \dots, K-1\}}$.
-

The penalization parameters γ_k are chosen using cross-validation, see Section 1 of Supplementary Material for precise statements about this procedure and about other numerical details.

3.2 Convergence to a stationary point

We are addressing here convergence properties of the QNEM algorithm described in Section 3.1 for the minimization of the objective function defined in equation (4). Let us denote

$$Q_n^{\text{pen}}(\theta, \theta^{(l)}) = Q_n(\theta, \theta^{(l)}) + \sum_{k=0}^{K-1} \gamma_k \left((1 - \eta) \|\beta_k\|_1 + \frac{\eta}{2} \|\beta_k\|_2^2 \right)$$

Convergence properties of the EM algorithm in a general setting are well known, see Wu.²³ In the QNEM algorithm, since we only improve $Q_n^{\text{pen}}(\theta, \theta^{(l)})$ instead of a minimization of $Q_n(\theta, \theta^{(l)})$, we are not in the EM algorithm setting but in a so called generalized EM (GEM) algorithm setting.²⁰ For such an algorithm, we do have the descent property, in the sense that the criterion function given in equation (4) is reduced at each iteration, namely

$$\ell_n^{\text{pen}}(\theta^{(l+1)}) \leq \ell_n^{\text{pen}}(\theta^{(l)})$$

Let us make two hypotheses.

Hypothesis 3: *The duration densities f_k are such that ℓ_n^{pen} is bounded for all θ .*

Hypothesis 4: *$Q_n^{\text{pen}}(\theta, \theta^{(l)})$ is continuous in θ and $\theta^{(l)}$, and for any fixed $\theta^{(l)}$, $Q_n^{\text{pen}}(\theta, \theta^{(l)})$ is a convex function in θ and is strictly convex in each coordinate of θ .*

Under Hypothesis 3, $l \mapsto \ell_n^{\text{pen}}(\theta^{(l)})$ decreases monotonically to some finite limit. By adding Hypothesis 4, convergence of the QNEM algorithm to a stationary point can be shown. In particular, the stationary point is here a local minimum.

Theorem 1: *Under Hypotheses 3 and 4, and considering the QNEM algorithm for the criterion function defined in equation (4), every cluster point $\bar{\theta}$ of the sequence $\{\theta^{(l)}; l = 0, 1, 2, \dots\}$ generated by the QNEM algorithm is a stationary point of the criterion function defined in equation (4).*

A proof is given in Section 2 of Supplementary Material.

3.3 Parameterization

Let us discuss here the parameterization choices we made in the experimental part. First, in many applications – including the one addressed in Section 5 – we are interested in identifying one subgroup of the population with a high risk of adverse event compared to the others. Then, in the following, we consider $Z \in \{0, 1\}$ where $Z = 1$ means high risk of early death and $Z = 0$ means low risk. Moreover, in such a setting where $K = 2$, one can compare the learned groups by the C-mix and the ones learned by the CURE model in terms of survival curves (see Figure 5).

To simplify notations and given the constraint formulated in equation (2), we set $\beta_0 = 0$ and we denote $\beta = \beta_1$ and $\pi_\beta(x)$ the conditional probability that a patient belongs to the group with high risk of death, given its covariates x .

In practice, we deal with discrete times in days. It turns out that the times of the data used for applications in Section 5 are well fitted by Weibull distributions. This choice of distribution is very popular in survival analysis, see for instance Klein and Moeschberger.¹⁸ We then first derive the QNEM algorithm with

$$f_k(t; \alpha_k) = (1 - \phi_k)^{t^{\mu_k}} - (1 - \phi_k)^{(t+1)^{\mu_k}}$$

with here $\alpha_k = (\phi_k, \mu_k) \in (0, 1) \times \mathbb{R}_+$, ϕ_k being the scale parameter and μ_k the shape parameter of the distribution.

As explained in the following Section 4, we select the best model using a cross-validation procedure based on the C-index metric, and the performances are evaluated according to both C-index and $\text{AUC}(t)$ metrics (see Sections 4.3 for details). Those two metrics have the following property: if we apply any mapping on the marker vector (predicted on a test set) such that the order between all vector coefficient values is conserved, then both C-index and $\text{AUC}(t)$ estimates remain unchanged. In other words, by denoting $(M_i)_{i \in \{1, \dots, n_{\text{test}}\}}$ the vector of markers predicted on a test set of n_{test} individuals, if ψ is a function such that for all $(i, j) \in \{1, \dots, n_{\text{test}}\}^2$, $(M_i < M_j \Rightarrow \psi(M_i) < \psi(M_j))$, then both C-index and $\text{AUC}(t)$ estimates induced by $(M_i)_{i \in \{1, \dots, n_{\text{test}}\}}$ or by $(\psi(M_i))_{i \in \{1, \dots, n_{\text{test}}\}}$ are the same.

The order in the marker coefficients is actually paramount when the performances are evaluated according to the mentioned metrics. Furthermore, it turns out that empirically, if we add a constraint on the mixture of Weibull that enforces an *order like* relation between the two distributions f_0 and f_1 , the performances are improved. To be more precise, the constraint to impose is that the two density curves do not intersect. We then choose to impose the following: the two scale parameters are equal, *i.e.* $\phi_0 = \phi_1 = \phi$. Indeed under this hypothesis, we do have that for all $\phi \in (0, 1)$, $(\mu_0 < \mu_1 \Rightarrow \forall t \in \mathbb{R}_+, f_0(t; \alpha_0) > f_1(t; \alpha_1))$.

With this Weibull parameterization, updates for α_k are not explicit in the QNEM algorithm, and consequently require some iterations of a minimization algorithm. Seeking to have explicit updates for α_k , we then derive the algorithm with geometric distributions instead of Weibull (geometric being a particular case of Weibull with $\mu_k = 1$), namely $f_k(t; \alpha_k) = \alpha_k(1 - \alpha_k)^{t-1}$ with $\alpha_k \in (0, 1)$.

With this parameterization, we obtain from equation (7)

$$\begin{aligned} \Lambda_{1,i}^{(l)} &= [\alpha_1^{(l)}(1 - \alpha_1^{(l)})^{y_i-1}]^{\delta_i} [(1 - \alpha_1^{(l)})^{y_i}]^{1-\delta_i} \pi_{\beta^{(l)}}(x_i) \quad \text{and} \\ \Lambda_{0,i}^{(l)} &= [\alpha_0^{(l)}(1 - \alpha_0^{(l)})^{y_i-1}]^{\delta_i} [(1 - \alpha_0^{(l)})^{y_i}]^{1-\delta_i} (1 - \pi_{\beta^{(l)}}(x_i)) \end{aligned}$$

which leads to the following explicit M-step

$$\alpha_0^{(l+1)} = \frac{\sum_{i=1}^n \delta_i (1 - q_i^{(l)})}{\sum_{i=1}^n (1 - q_i^{(l)}) y_i} \quad \text{and} \quad \alpha_1^{(l+1)} = \frac{\sum_{i=1}^n \delta_i q_i^{(l)}}{\sum_{i=1}^n q_i^{(l)} y_i}$$

In this setting, implementation is hence straightforward. Note that Hypotheses 3 and 4 are immediately satisfied with this geometric parameterization.

In Section 5, we note that performances are similar for the C-mix model with Weibull or geometric distributions on all considered biomedical datasets. The geometric parameterization leading to more straightforward computations, it is the one used to parameterize the C-mix model in what follows, if not otherwise stated. Let us focus now on the performance evaluation of the C-mix model and its comparison with the Cox PH and CURE models, both regularized with the Elastic-Net.

4 Performance evaluation

In this section, we first briefly introduce the models we consider for performance comparisons. Then, we provide details regarding the simulation study and data generation. The chosen metrics for evaluating performances are then presented, followed by the results.

4.1 Competing models

The first model we consider is the Cox PH model penalized by the Elastic-Net, denoted Cox PH in the following. In this model introduced in Cox,¹² the partial log-likelihood is given by

$$\ell_n^{\text{cox}}(\beta) = n^{-1} \sum_{i=1}^n \delta_i \left(x_i^\top \beta - \log \sum_{i': y_{i'} \geq y_i} \exp(x_{i'}^\top \beta) \right)$$

We use respectively the R packages `survival` and `glmnet`¹³ for the partial log-likelihood and the minimization of the following quantity

$$-\ell_n^{\text{cox}}(\beta) + \gamma \left((1 - \eta) \|\beta\|_1 + \frac{\eta}{2} \|\beta\|_2^2 \right)$$

where γ is chosen by the same cross-validation procedure than the C-mix model, for a given η (see Section 1 of Supplementary Material). Ties are handled through the Breslow approximation of the partial likelihood.²⁴

We remark that the model introduced in this paper cannot be reduced to a Cox model. Indeed, the C-mix model intensity can be written (in the geometric case)

$$\lambda(t) = \frac{\alpha_1(1 - \alpha_1)^{t-1} + \alpha_0(1 - \alpha_0)^{t-1} \exp(x^\top \beta)}{(1 - \alpha_1)^t + (1 - \alpha_0)^t \exp(x^\top \beta)}$$

while it is given by equation (1) in the Cox model.

Finally, we consider the CURE¹⁴ model penalized by the Elastic-Net and denoted CURE in the following, with a logistic function for the incidence part and a parametric survival model for $S(t|Z = 1)$, where $Z = 0$ means that patient is cured, $Z = 1$ means that patient is not cured, and $S(t) = \exp(-\int_0^t \lambda(s) ds)$ denotes the survival function. In this model, we then have $S(t|Z = 0)$ constant and equal to 1. We add an Elastic-Net regularization term, and since we were not able to find any open source package where CURE models were implemented with a regularized objective, we used the QNEM algorithm in the particular case of CURE model. We just add the constraint that the geometric distribution $\mathcal{G}(\alpha_0)$ corresponding to the cured group of patients ($Z = 0$) has a parameter $\alpha_0 = 0$, which does not change over the algorithm iterations. The QNEM algorithm can be used in this particular case, were some terms have disappeared from the completed log-likelihood, since in the CURE model case we have $\{i \in \{1, \dots, n\} : z_i = 0, \delta_i = 1\} = \emptyset$. Note that in the original introduction of the CURE model in Farewell,¹⁴ the density of uncured patients directly depends on individual patient covariates, which is not the case here.

We also give additional simulation settings in Section 3 of Supplementary Material. First, the case where $d \gg n$, including a comparison of the screening strategy we use in Section 5 with the iterative sure independence screening²⁵ (ISIS) method. We also add simulations where data are generated according to the C-mix model with gamma distributions instead of geometric ones, and include the accelerated failure time model²⁶ (AFT) in the performances comparison study.

4.2 Simulation design

In order to assess the proposed method, we perform an extensive Monte Carlo simulation study. Since we want to compare the performances of the three models mentioned above, we consider three simulation cases for the time

distribution: one for each competing model. We first choose a coefficient vector $\beta = (\underbrace{v, \dots, v}_s, 0, \dots, 0) \in \mathbb{R}^d$, with $v \in \mathbb{R}$ being the value of the active coefficients and $s \in \{1, \dots, d\}$ a sparsity parameter. For a desired low-risk patients proportion $\pi_0 \in [0, 1]$, the high-risk patients index set is given by

$$\mathcal{H} = \{\lfloor (1 - \pi_0) \times n \rfloor \text{ random samples without replacement}\} \subset \{1, \dots, n\},$$

where $\lfloor a \rfloor$ denotes the largest integer less than or equal to $a \in \mathbb{R}$. For the generation of the covariates matrix, we first take $[x_{ij}] \in \mathbb{R}^{n \times d} \sim \mathcal{N}(0, \Sigma(\rho))$, with $\Sigma(\rho)$ a $(d \times d)$ Toeplitz covariance matrix²⁷ with correlation $\rho \in (0, 1)$. We then add a gap $\in \mathbb{R}^+$ value for patients $i \in \mathcal{H}$ and subtract it for patients $i \notin \mathcal{H}$, only on active covariates plus a proportion $r_{cf} \in [0, 1]$ of the non-active covariates considered as confusion factors, that is

$$x_{ij} \leftarrow x_{ij} \pm \text{gap for } j \in \{1, \dots, s, \dots, \lfloor (d - s)r_{cf} \rfloor\}$$

Note that this is equivalent to generate the covariates according to a gaussian mixture.

Then we generate $Z_i \sim \mathcal{B}(\pi_{\beta}(x_i))$ in the C-mix or CURE simulation case, where $\pi_{\beta}(x_i)$ is computed given equation (3), with geometric distributions for the durations (see Section 3.3). We obtain $T_i \sim \mathcal{G}(\alpha_{Z_i})$ in the C-mix case, and $T_i \sim \infty \mathbb{1}_{\{Z_i=0\}} + \mathcal{G}(\alpha_1) \mathbb{1}_{\{Z_i=1\}}$ in the CURE case. For the Cox PH model, we take $T_i \sim -\log(U_i) \exp(-x_i^\top \beta)$, with $U_i \sim \mathcal{U}([0, 1])$ and where $\mathcal{U}([a, b])$ stands for the uniform distribution on a segment $[a, b]$.

The distribution of the censoring variable C_i is geometric $\mathcal{G}(\alpha_c)$, with $\alpha_c \in (0, 1)$. The parameter α_c is tuned to maintain a desired censoring rate $r_c \in [0, 1]$, using a formula given in Section 4 of Supplementary Material. The values of the chosen hyper parameters are summarized in Table 1.

Note that when simulating under the CURE model, the proportion of censored time events is at least equal to π_0 : we then choose $\pi_0 = 0.2$ for the CURE simulations only.

Finally, we want to assess the stability of the C-mix model in terms of variable selection and compare it to the CURE and Cox PH models. To this end, we follow the same simulation procedure explained in the previous lines. For each simulation case, we make vary the two hyper-parameters that impact the most the stability of the variable selection, that is the gap varying in $[0, 2]$ and the confusion rate r_{cf} varying in $[0, 1]$. All other hyper-parameters are the same than in Table 1, except $s = 150$ and with the choice $(n, d) = (200, 300)$. For a given hyper-parameters configuration (gap, r_{cf}), we use the following approach to evaluate the variable selection power of the models. Denoting $\hat{\beta}_i = |\hat{\beta}_i| / \max\{|\hat{\beta}_i|, i \in \{1, \dots, d\}\}$, if we consider that $\hat{\beta}_i$ is the predicted probability that the true β_i equals v , then we are in a binary prediction setting and we use the resulting AUC of this problem. Explicit examples of such AUC computations are given in Section 5 of Supplementary Material.

4.3 Metrics

We detail in this section the metrics considered to evaluate risk prediction performances. Let us denote by M the marker under study. Note that $M = \pi_{\hat{\beta}}(X)$ in the C-mix and the CURE model cases, and $M = \exp(X^\top \hat{\beta}^{\text{cox}})$ in the Cox PH model case. We denote by h the probability density function of marker M , and assume that the marker is measured once at $t = 0$.

For any threshold ξ , cumulative true positive rates and dynamic false positive rates are two functions of time respectively defined as $\text{TPR}^{\text{C}}(\xi, t) = \mathbb{P}[M > \xi | T \leq t]$ and $\text{FPR}^{\text{D}}(\xi, t) = \mathbb{P}[M > \xi | T > t]$. Then, as introduced in Heagerty et al.,²⁸ the cumulative dynamic time-dependent AUC is defined as follows

$$\text{AUC}^{\text{C, D}}(t) = \int_{-\infty}^{\infty} \text{TPR}^{\text{C}}(\xi, t) \left| \frac{\partial \text{FPR}^{\text{D}}(\xi, t)}{\partial \xi} \right| d\xi$$

Table 1. Hyper-parameters choice for simulation.

η	n	d	s	r_{cf}	v	ρ	π_0	gap	r_c	α_0	α_1
0.1	100, 200, 500	30, 100	10	0.3	1	0.5	0.75	0.1, 0.3, 1	0.2, 0.5	0.01	0.5

that we simply denote $AUC(t)$ in the following. We use the Inverse Probability of Censoring Weighting (IPCW) estimate of this quantity with a Kaplan–Meier estimator of the conditional survival function $\mathbb{P}[T > t | M = m]$, as proposed in Blanche et al.²⁹ and already implemented in the R package `timeROC`.

A common concordance measure that does not depend on time is the C-index³⁰ defined by

$$C = \mathbb{P}[M_i > M_j | T_i < T_j]$$

with $i \neq j$ two independent patients (which 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 C_τ defined by

$$C_\tau = \mathbb{P}[M_i > M_j | Y_i < Y_j, Y_i < \tau]$$

with τ corresponding to the fixed and prespecified follow-up period duration.³¹ A Kaplan–Meier estimator for the censoring distribution leads to a nonparametric and consistent estimator of C_τ ,³² already implemented in the R package `survival`.

Hence in the following, we consider both $AUC(t)$ and C-index metrics to assess performances.

4.4 Results of simulation

We present now the simulation results concerning the C-index metric in the case $(d, r_c) = (30, 0.5)$ in Table 2. See Section 6 of Supplementary Material for results on other configurations for (d, r_c) . Each value is obtained by computing the C-index average and standard deviation (in parenthesis) over 100 simulations. The $AUC(t)$ average (bold line) and standard deviation (bands) over the same 100 simulations are then given in Figure 1, where $n = 100$. Note that the value of the gap can be viewed as a difficulty level of the problem, since the higher the value of the gap, the clearer the separation between the two populations (low-risk and high-risk patients).

The results measured both by $AUC(t)$ and C-index lead to the same conclusion: the C-mix model almost always leads to the best results, even under model misspecification, i.e. when data are generated according to the CURE or Cox PH model. Namely, under CURE simulations, C-mix and CURE give very close results, with a strong improvement over Cox PH. Under Cox PH and C-mix simulations, C-mix outperforms both Cox PH and CURE. Surprisingly enough, this exhibits a strong generalization property of the C-mix model, over both Cox PH and CURE. Note that this phenomenon is particularly strong for small gap values, while with an increasing gap (or an increasing sample size n), all procedures barely exhibit the same performance. It can be first explained by the nonparametric baseline function in the Cox PH model, and second by the fact that unlike the Cox PH model, the C-mix and CURE models exploit directly the mixture aspect.

Finally, Figure 2 gives the results concerning the stability of the variable selection aspect of the competing models. The C-mix model appears to be the best method as well considering the variable selection aspect, even under model misspecification. We notice a general behaviour of our method that we describe in the following, which is also shared by the CURE model only when the data are simulated according to itself, and which justifies the log scale for the gap to clearly distinguish the three following phases. For very small gap values (less than 0.2), the confusion rate r_{cf} does not impact the variable selection performances, since adding very small gap values to the covariates is almost imperceptible. This means that the resulting AUC is the same when there is no confusion factors and when $r_{cf} = 1$ (that is when there are half active covariates and half confusion ones). For medium gap values (saying between 0.2 and 1), the confusion factors are more difficult to identify by the model as there number goes up (that is when r_{cf} increases), which is precisely the confusion factors effect we expect to observe. Then, for large gap values (more than 1), the model succeeds in vanishing properly all confusion factors since the two subpopulations are more clearly separated regarding the covariates, and the problem becomes naturally easier as the gap increases.

5 Application to genetic data

In this section, we apply our method on three genetic datasets and compare its performance to the Cox PH and CURE models. We extracted normalized expression data and survival times Y in days from breast invasive carcinoma (BRCA, $n = 1211$), glioblastoma multiforme (GBM, $n = 168$) and kidney renal clear cell carcinoma (KIRC, $n = 605$).

Table 2. Average C-index on 100 simulated data and standard deviation in parenthesis, with $d = 30$ and $r_c = 0.5$.

Simulation	Gap	Estimation								
		$n = 100$			$n = 200$			$n = 500$		
		C-mix	CURE	Cox PH	C-mix	CURE	Cox PH	C-mix	CURE	Cox PH
C-mix	0.1	0.786 (0.057)	0.745 (0.076)	0.701 (0.075)	0.792 (0.040)	0.770 (0.048)	0.739 (0.055)	0.806 (0.021)	0.798 (0.023)	0.790 (0.024)
	0.3	0.796 (0.055)	0.739 (0.094)	0.714 (0.088)	0.794 (0.036)	0.760 (0.058)	0.744 (0.055)	0.801 (0.021)	0.784 (0.027)	0.783 (0.026)
	I	0.768 (0.062)	0.734 (0.084)	0.756 (0.066)	0.766 (0.043)	0.736 (0.054)	0.764 (0.042)	0.772 (0.026)	0.761 (0.027)	0.772 (0.025)
CURE	0.1	0.770 (0.064)	0.772 (0.062)	0.722 (0.073)	0.790 (0.038)	0.790 (0.038)	0.758 (0.049)	0.798 (0.025)	0.799 (0.024)	0.787 (0.025)
	0.3	0.733 (0.073)	0.732 (0.072)	0.686 (0.072)	0.740 (0.053)	0.741 (0.053)	0.714 (0.060)	0.751 (0.029)	0.751 (0.029)	0.738 (0.030)
	I	0.659 (0.078)	0.658 (0.078)	0.635 (0.070)	0.658 (0.053)	0.658 (0.053)	0.647 (0.047)	0.657 (0.031)	0.657 (0.031)	0.656 (0.032)
Cox PH	0.1	0.940 (0.041)	0.937 (0.044)	0.850 (0.097)	0.959 (0.021)	0.958 (0.020)	0.915 (0.042)	0.964 (0.012)	0.964 (0.012)	0.950 (0.016)
	0.3	0.956 (0.030)	0.955 (0.029)	0.864 (0.090)	0.966 (0.020)	0.965 (0.020)	0.926 (0.043)	0.968 (0.013)	0.969 (0.012)	0.959 (0.016)
	I	0.983 (0.016)	0.985 (0.015)	0.981 (0.019)	0.984 (0.012)	0.985 (0.011)	0.988 (0.010)	0.984 (0.007)	0.985 (0.006)	0.990 (0.005)

For each configuration, the best result appears in bold.

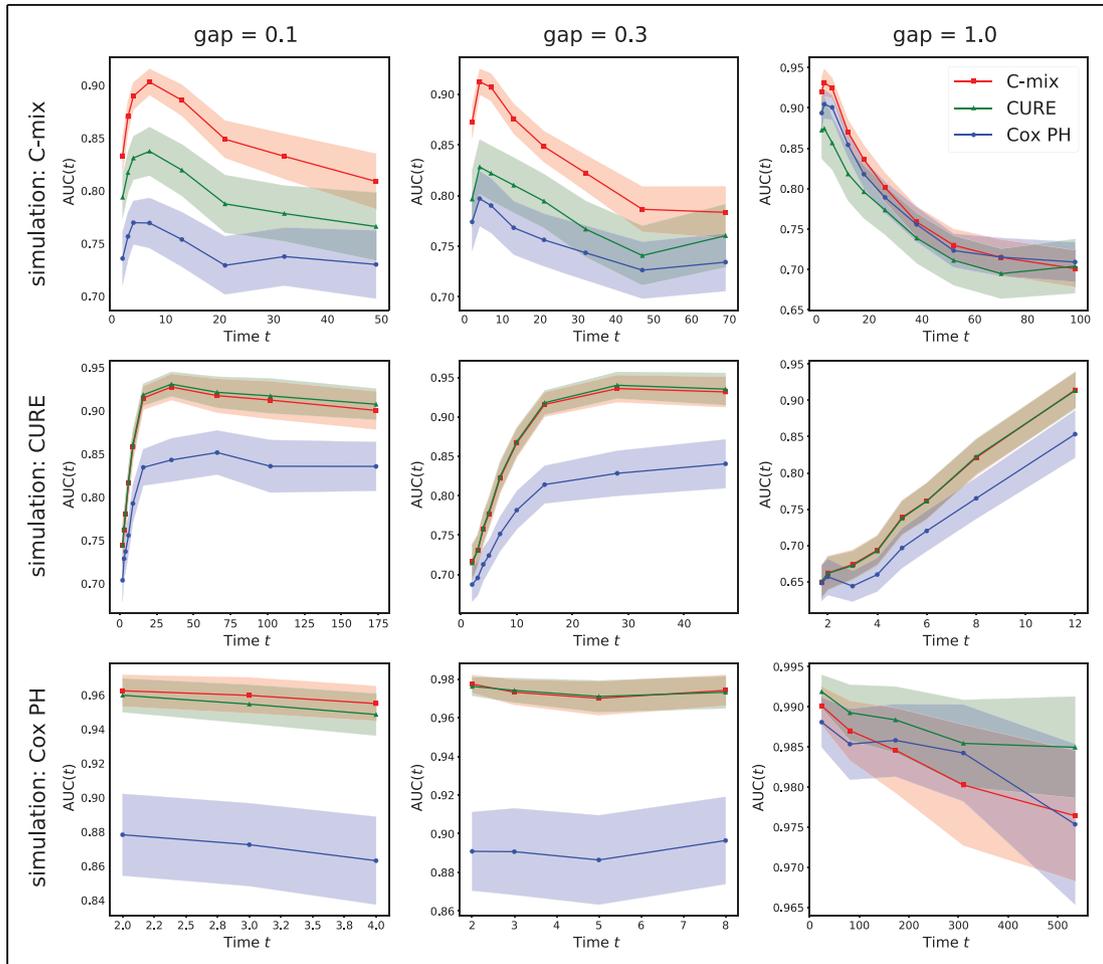


Figure 1. Average (bold lines) and standard deviation (bands) for $AUC(t)$ on 100 simulated data with $n = 100$, $d = 30$ and $r_c = 0.5$. Rows correspond to the model simulated (cf. Section 4.2) while columns correspond to different gap values (the problem becomes more difficult as the gap value decreases). Surprisingly, our method gives almost always the best results, even under model misspecification (see Cox PH and CURE simulation cases on the second and third rows).

These datasets are available on The Cancer Genome Atlas (TCGA) platform, which aims at accelerating the understanding of the molecular basis of cancer through the application of genomic technologies, including large-scale genome sequencing. For each patient, 20,531 covariates corresponding to the normalized gene expressions are available. We randomly split all datasets into a training set and a test set (30% for testing, cross-validation is done on the training).

We compare the three models both in terms of C-index and $AUC(t)$ on the test sets. Inference of the Cox PH model fails in very high dimension on the considered data with the `glmnet` package. We therefore make a first variable selection (screening) among the 20,531 covariates. To do so, we compute the C-index obtained by univariate Cox PH models (not to confer advantage to our method), namely Cox PH models fitted on each covariate separately. We then ordered the obtained 20,531 C-indexes by decreasing order and extracted the top $d = 100$, $d = 300$ and $d = 1000$ covariates. We then apply the three methods on the obtained covariates.

The results in terms of $AUC(t)$ curves are given in Figure 3 for $d = 300$, where we distinguish the C-mix model with geometric or Weibull distributions.

Then it appears that the performances are very close in terms of $AUC(t)$ between the C-mix model with geometric or Weibull distributions, which is also validated if we compare the corresponding C-index for these two parameterizations in Table 3.

Similar conclusions in terms of C-index, $AUC(t)$ and computing time can be made on all considered datasets and for any choice of d . Hence, as already mentioned in Section 3.3, we only concentrate on the geometric parameterization for the C-mix model. The results in terms of C-index are then given in Table 4.

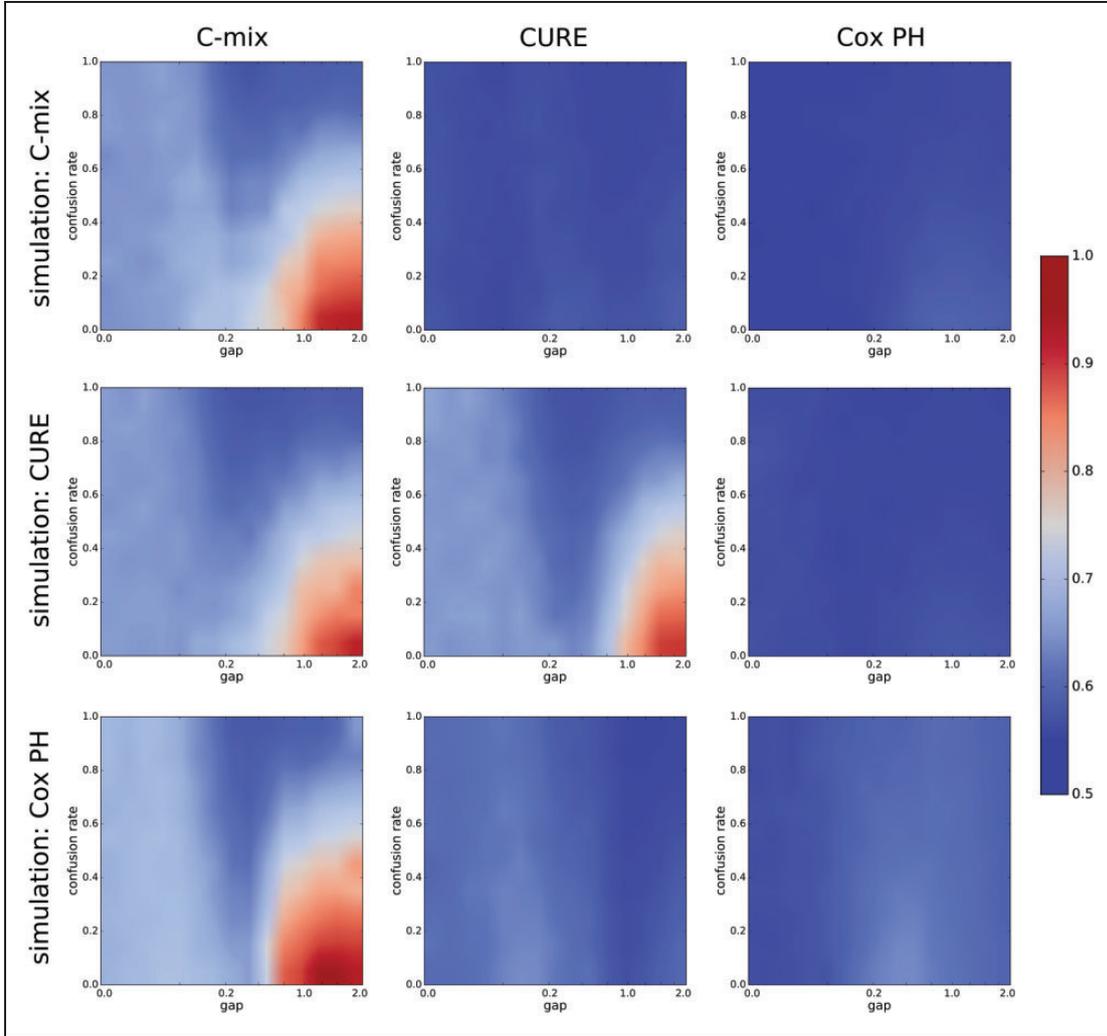


Figure 2. Average AUC calculated according to Section 4.2 and obtained after 100 simulated data for each (gap, r_{ct}) configuration (a grid of 20×20 different configurations is considered). A gaussian interpolation is then performed to obtain smooth figures. Note that the gap values are *log*-scaled. Rows correspond to the model simulated while columns correspond to the model under consideration for the variable selection evaluation procedure. Our method gives the best results in terms of variable selection, even under model misspecification.

A more direct approach to compare performances between models, rather than only focus on the marker order aspect, is to predict the survival of patients in the test set within a specified short time. For the Cox PH model, the survival $\mathbb{P}[T_i > t | X_i = x_i]$ for patient i in the test set is estimated by

$$\hat{S}_i(t | X_i = x_i) = [\hat{S}_0^{\text{cox}}(t)]^{\exp(x_i^\top \hat{\beta}^{\text{cox}})}$$

where \hat{S}_0^{cox} is the estimated survival function of baseline population ($x = 0$) obtained using the Breslow estimate of λ_0 .²⁴ For the CURE or the C-mix models, it is naturally estimated by

$$\hat{S}_i(t | X_i = x_i) = \pi_{\hat{\beta}}(x_i) \hat{S}_1(t) + (1 - \pi_{\hat{\beta}}(x_i)) \hat{S}_0(t)$$

where \hat{S}_0 and \hat{S}_1 are the Kaplan–Meier estimators³³ of the low- and high-risk subgroups respectively, learned by the C-mix or CURE models (patients with $\pi_{\hat{\beta}}(x_i) > 0.5$ are clustered in the high-risk subgroup, others in the low-

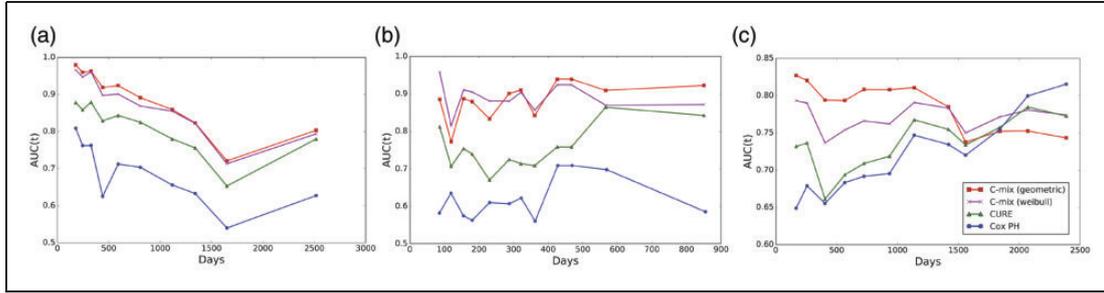


Figure 3. AUC(t) comparison on the three TCGA datasets considered, for $d=300$. We observe that C-mix model leads to the best results (higher is better) and outperforms both Cox PH and CURE in all cases. Results are similar in terms of performances for the C-mix model with geometric or Weibull distributions. (a) BRCA; (b) GBM; (c) KIRC.

Table 3. C-index comparison between geometric or Weibull parameterizations for the C-mix model on the three TCGA datasets considered (with $d=300$).

Parameterization		Geometric	Weibull
Cancer	BRCA	0.782	0.780
	GBM	0.755	0.754
	KIRC	0.849	0.835

In all cases, results are very similar for the two distribution choices.

Table 4. C-index comparison on the three TCGA datasets considered.

Cancer Model	d	BRCA			GBM			KIRC		
		C-mix	CURE	Cox PH	C-mix	CURE	Cox PH	C-mix	CURE	Cox PH
	100	0.792	0.764	0.705	0.826	0.695	0.571	0.768	0.732	0.716
	300	0.782	0.753	0.723	0.849	0.697	0.571	0.755	0.691	0.698
	1000	0.817	0.613	0.577	0.775	0.699	0.592	0.743	0.690	0.685

In all cases, C-mix gives the best results (in bold).

risk one). The corresponding estimated survival curves are given in Figure 4. We observe that the subgroups obtained by the C-mix are more clearly separated in terms of survival than those obtained by the CURE model.

For a given time ϵ , one can now use $\hat{S}_i(\epsilon|X_i = x_i)$ for each model to predict whether or not $T_i > \epsilon$ on the test set, resulting on a binary classification problem that we assess using the classical AUC score. By moving ϵ within the first years of follow-up, since it is the more interesting for physicians in practice, one obtains the curves given in Figure 5.

Let us now focus on the runtime comparison between the models in Table 5. We choose the BRCA dataset to illustrate this point, since it is the larger one ($n=1211$) and consequently provides more clearer time-consuming differences.

We also notice that despite using the same QNEM algorithm steps, our CURE model implementation is slower since convergence takes more time to be reached, as shows Figure 6.

In Section 7 of Supplementary Material, the top 20 selected genes for each cancer type and for all models are presented (for $d=300$). Literature on those genes is mined to estimate two simple scores that provide information about how related they are to cancer in general first, and second to cancer plus the survival aspect, according to scientific publications. It turns out that some genes have been widely studied in the literature (e.g. FLT3 for the GBM cancer), while for others, very few publications were retrieved (e.g. TRMT2B still for the GBM cancer).

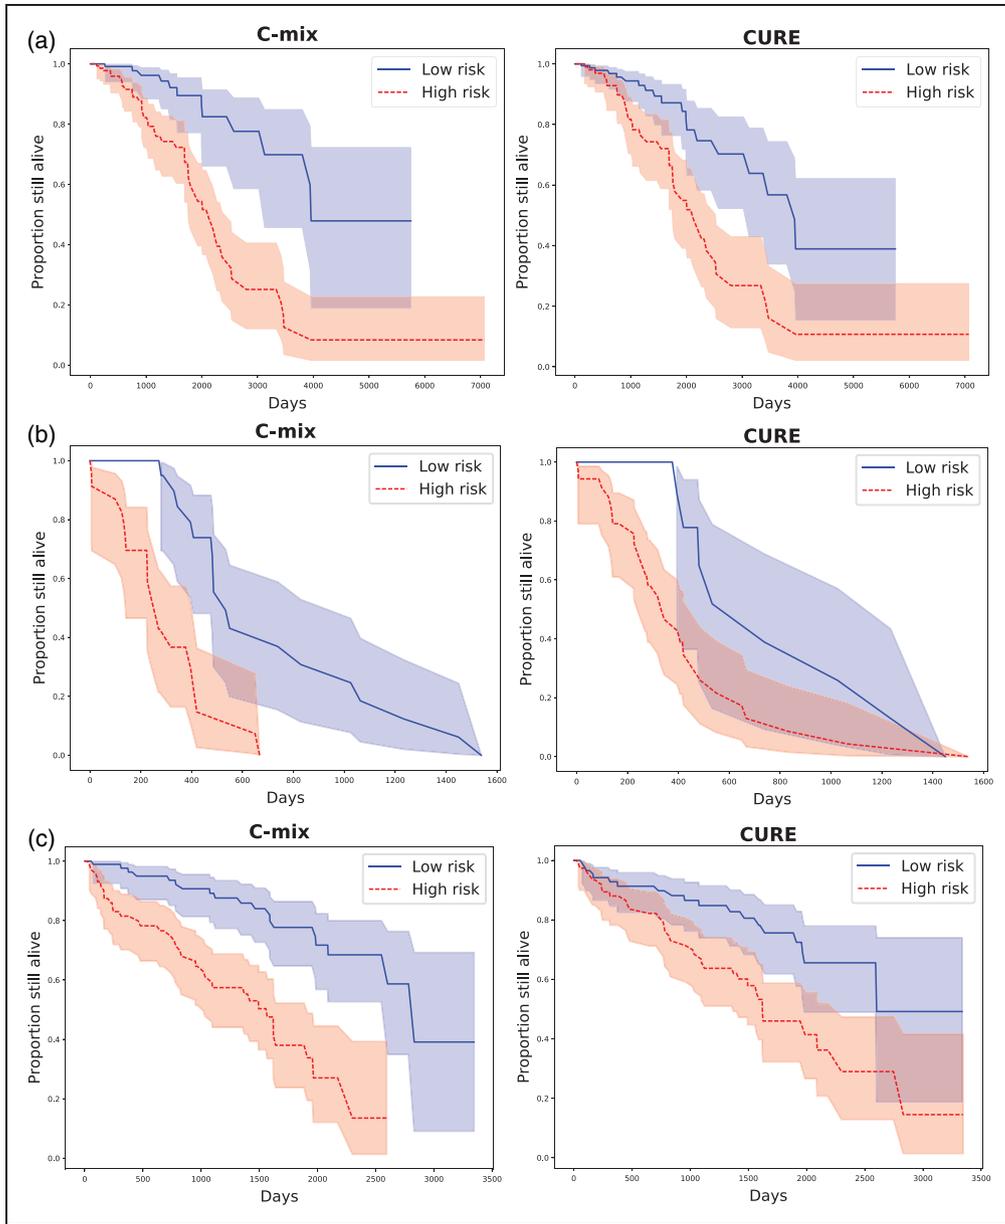


Figure 4. Estimated survival curves per subgroups (blue for low risk and red for high risk) with the corresponding 95% confidence bands for the C-mix and CURE models: BRCA in column (a), GBM in column (b) and KIRC in column (c).

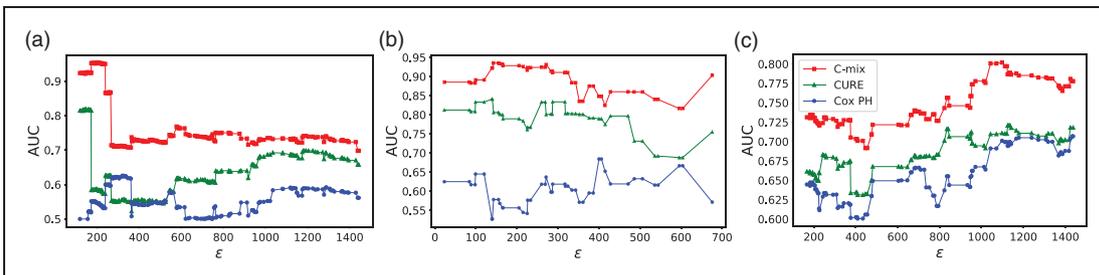


Figure 5. Comparison of the survival prediction performances between models on the three TCGA datasets considered (still with $d=300$). Performances are, once again, much better for the C-mix over the two other standard methods. (a) BRCA; (b) GBM; (c) KIRC.

Table 5. Computing time comparison in second on the BRCA dataset ($n = 1211$), with corresponding C-index in parenthesis and best result in bold in each case.

Model		C-mix	CURE	Cox PH
d	100	0.025 (0.792)	1.992 (0.764)	0.446 (0.705)
	300	0.027 (0.782)	2.343 (0.753)	0.810 (0.723)
	1000	0.139 (0.817)	12.067 (0.613)	2.145 (0.577)

This times concern the learning task for each model with the best hyper parameter selected after the cross validation procedure. It turns out that our method is by far the fastest in addition to providing the best performances. In particular, the QNEM algorithm is faster than the R implementation `glmnet`.

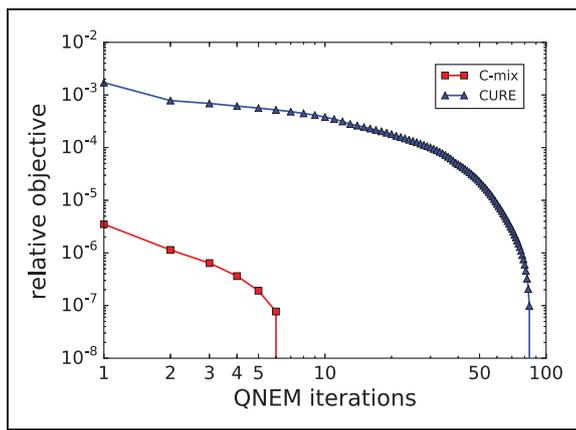


Figure 6. Convergence comparison between C-mix and CURE models through the QNEM algorithm. The relative objective is here defined at iteration l as $(\ell_n^{\text{pen}}(\hat{\theta}^{(l)}) - \ell_n^{\text{pen}}(\hat{\theta})) / \ell_n^{\text{pen}}(\hat{\theta})$, where $\hat{\theta}$ is naturally the parameter vector returned at the end of the QNEM algorithm, that is once convergence is reached. Note that both iteration and relative objective axis are \log -scaled for clarity. We observe that convergence for the C-mix model is dramatically faster than the CURE one.

6 Concluding remarks

In this paper, a mixture model for censored durations (C-mix) has been introduced, and a new efficient estimation algorithm (QNEM) has been derived, that considers a penalization of the likelihood in order to perform covariate selection and to prevent overfitting. A strong improvement is provided over the CURE and Cox PH approaches (both penalized by the Elastic-Net), which are, by far, the most widely used for biomedical data analysis. But more importantly, our method detects relevant subgroups of patients regarding their risk in a supervised learning procedure, and takes advantage of this identification to improve survival prediction over more standard methods. An extensive Monte Carlo simulation study has been carried out to evaluate the performance of the developed estimation procedure. It showed that our approach is robust to model misspecification. The proposed methodology has then been applied on three high-dimensional datasets. On these datasets, C-mix outperforms both Cox PH and CURE, in terms of $\text{AUC}(t)$, C-index or survival prediction. Moreover, many gene expressions pinpointed by the feature selection aspect of our regularized method are relevant for medical interpretations (e.g. NFKBIA, LEF1, SUSD3 or FAIM3 for the BRCA cancer, see Zhou et al.³⁴ or Oskarsson et al.³⁵), whilst others must involve further investigations in the genetic research community. Finally, our analysis provides, as a by-product, a new robust implementation of CURE models in high dimension.

Software

All the methodology discussed in this paper is implemented in Python. The code is available from <https://github.com/SimonBussy/C-mix> in the form of annotated programs, together with a notebook tutorial.

Acknowledgements

The results shown in this paper are based on data generated by the TCGA Research Network and freely available from <http://cancergenome.nih.gov/>

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Supplementary material

Supplementary material available at *Statistical Methods in Medical Research* online contains technical sections with numerical details, a proof of Theorem 1, additional simulations and details about the selected genes per models on the TCGA datasets.

ORCID iD

Simon Bussy  <http://orcid.org/0000-0002-6059-0685>

References

1. Alizadeh AA, Eisen MB, Davis RE, et al. Distinct types of diffuse large b-cell lymphoma identified by gene expression profiling. *Nature* 2000; **403**: 503–511.
2. Rosenwald A, Wright G, Chan WC, et al. The use of molecular profiling to predict survival after chemotherapy for diffuse large-b-cell lymphoma. *New Engl J Med* 2002; **346**: 1937–1947.
3. Golub TR, Slonim DK, Tamayo P, et al. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* 1999; **286**: 531–537.
4. Hastie T, Tibshirani R, Botstein D, et al. Supervised harvesting of expression trees. *Genome Biol* 2001; **2**: research0003.1–research0003.12.
5. Tibshirani R, Hastie T, Narasimhan B, et al. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proc Natl Acad Sci* 2002; **99**: 6567–6572.
6. Beer DG, Kardia SL, Huang CC, et al. Gene-expression profiles predict survival of patients with lung adenocarcinoma. *Nat Med* 2002; **8**: 816–824.
7. Bhattacharjee A, Richards WG, Staunton J, et al. Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *Proc Natl Acad Sci* 2001; **98**: 13790–13795.
8. Sørlie T, Perou CM, Tibshirani R, et al. Gene expression patterns of breast carcinomas distinguish tumor subclasses with clinical implications. *Proc Natl Acad Sci* 2001; **98**: 10869–10874.
9. Shipp MA, Ross KN, Tamayo P, et al. Diffuse large b-cell lymphoma outcome prediction by gene-expression profiling and supervised machine learning. *Nat Med* 2002; **8**: 68–74.
10. Van't Veer LJ, Dai H, Van De Vijver MJ, et al. Gene expression profiling predicts clinical outcome of breast cancer. *Nature* 2002; **415**: 530–536.
11. Banfield JD and Raftery AE. Model-based gaussian and non-gaussian clustering. *Biometrics* 1993; **49**: 803–821.
12. Cox DR. Regression models and life-tables. *J Roy Stat Soc B* 1972; **34**: 187–220.
13. Simon N, Friedman J, Hastie T, et al. Regularization paths for Coxs proportional hazards model via coordinate descent. *J Stat Softw* 2011; **39**: 1–13.
14. Farewell V. The use of mixture models for the analysis of survival data with long-term survivors. *Biometrics* 1982; **38**: 1041–1046.
15. Kuk AY and Chen CH. A mixture model combining logistic regression with proportional hazards regression. *Biometrika* 1992; **79**: 531–541.
16. Kuo L and Peng F. A mixture-model approach to the analysis of survival data. *Biosstat – Basel* 2000; **5**: 255–272.
17. De Angelis R, Capocaccia R, Hakulinen T, et al. Mixture models for cancer survival analysis: application to population-based data with covariates. *Stat Med* 1999; **18**: 441–454.
18. Klein JP and Moeschberger ML. *Survival analysis: techniques for censored and truncated data*. New York: Springer Science & Business Media, 2005.

19. Zou H and Hastie T. Regularization and variable selection via the elastic net. *J R Stat Soc Series B Stat Methodol* 2005; **67**: 301–320.
20. Dempster A, Laird N and Rubin D. Maximum likelihood from incomplete data via the EM algorithm. *J Roy Stat Soc B* 1977; **39**: 1–38.
21. Zhu C, Byrd RH, Lu P, et al. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans Math Soft* 1997; **23**: 550–560.
22. Andrew G and Gao J. Scalable training of L1-regularized log-linear models. In: *International conference on machine learning*. New York: ACM, 2007, pp. 33–40.
23. Wu CJ. On the convergence properties of the EM algorithm. *Ann Stat* 1983; **11**: 95–103.
24. Breslow NE. Contribution to the discussion of the paper by DR Cox. *J Roy Stat Soc B* 1972; **34**: 216–217.
25. Fan J, Feng Y, Wu Y, et al. High-dimensional variable selection for Coxs proportional hazards model. In: *Borrowing strength: theory powering applications—a festschrift for Lawrence D. Brown*. Institute of Mathematical Statistics, Beachwood, OH, 2010, pp. 70–86.
26. Wei LJ. The accelerated failure time model: a useful alternative to the cox regression model in survival analysis. *Stat Med* 1992; **11**: 1871–1879.
27. Mukherjee BN and Maiti SS. On some properties of positive definite Toeplitz matrices and their possible applications. *Linear Algebra Appl* 1988; **102**: 211–240.
28. Heagerty PJ, Lumley T and Pepe MS. Time-dependent roc curves for censored survival data and a diagnostic marker. *Biometrics* 2000; **56**: 337–344.
29. Blanche P, Latouche A and Viallon V. Time-dependent AUC with right-censored data: a survey. *Risk Assess Eval Pred* 2013; **215**: 239–251.
30. Harrell FE, Lee KL and Mark DB. Tutorial in biostatistics multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Stat Med* 1996; **15**: 361–387.
31. Heagerty PJ and Zheng Y. Survival model predictive accuracy and roc curves. *Biometrics* 2005; **61**: 92–105.
32. Uno H, Cai T, Pencina MJ, et al. On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Stat Med* 2011; **30**: 1105–1117.
33. Kaplan EL and Meier P. Nonparametric estimation from incomplete observations. *J Am Stat Assoc* 1958; **53**: 457–481.
34. Zhou Y, Yau C, Gray JW, et al. Enhanced NF κ B and AP-1 transcriptional activity associated with antiestrogen resistant breast cancer. *BMC Cancer* 2007; **7**: 1.
35. Oskarsson T, Acharyya S, Zhang XH, et al. Breast cancer cells produce tenascin c as a metastatic niche component to colonize the lungs. *Nat Med* 2011; **17**: 867–874.