

Vivien Goepp*, Jean-Christophe Thalabard, Grégory Nuel and Olivier Bouaziz

Regularized bidimensional estimation of the hazard rate

<https://doi.org/10.1515/ijb-2019-0003>

Received January 7, 2019; accepted February 26, 2021; published online March 24, 2021

Abstract: In epidemiological or demographic studies, with variable age at onset, a typical quantity of interest is the incidence of a disease (for example the cancer incidence). In these studies, the individuals are usually highly heterogeneous in terms of dates of birth (the cohort) and with respect to the calendar time (the period) and appropriate estimation methods are needed. In this article a new estimation method is presented which extends classical age-period-cohort analysis by allowing interactions between age, period and cohort effects. We introduce a bidimensional regularized estimate of the hazard rate where a penalty is introduced on the likelihood of the model. This penalty can be designed either to smooth the hazard rate or to enforce consecutive values of the hazard to be equal, leading to a parsimonious representation of the hazard rate. In the latter case, we make use of an iterative penalized likelihood scheme to approximate the L_0 norm, which makes the computation tractable. The method is evaluated on simulated data and applied on breast cancer survival data from the SEER program.

Keywords: adaptive ridge procedure; age-period-cohort analysis; penalized likelihood; piecewise constant hazard; survival analysis.

1 Introduction

In epidemiological or demographic studies, with variable age at onset, a typical quantity of interest is the incidence or the hazard rate of a disease (for example the cancer incidence). In these studies, individuals are recruited and followed-up during a long period of time, usually from birth. The data are then reported either in the form of registers, which contain the number of observed cases and the number of individuals at risk to contract the disease, or in the form of the observed time for each individual. These types of studies are of great interest for the epidemiologist, especially when the event of interest will tend to occur at late ages, such as in cancer studies. However, these data are usually highly heterogeneous in terms of dates of birth and with respect to the calendar time. In such cases, it is therefore very important to take into account the variability of the age, the cohort (date of birth) and the period (the calendar time) in the hazard rate estimation. This is usually done using age-period-cohort estimation methods (see Yang and Land [1], and citations therein).

In age-period-cohort analysis, the effects of age, period and cohort are fit as factor variables in a regression model where the output is the logarithm of the hazard rate. However, this induces an identifiability problem due to the relationship: period = age + cohort. There have been several solutions proposed to this problem. Osmond and Gardner [2] proposed to compute each submodel (age-cohort, age-period, and period-cohort)

*Corresponding author: Vivien Goepp, MAP5, CNRS UMR 8145, 45, rue des Saints-Pères, 75006, Paris, France; MINES ParisTech, CBIO–Centre for Computational Biology, PSL Research University, 75006, Paris, France; Institut Curie, PSL Research University, 75005, Paris, France; and Inserm, U900, Paris, France, E-mail: vivien.goepp@mines-paristech.fr. <https://orcid.org/0000-0001-6961-4260>

Jean-Christophe Thalabard, MAP5, CNRS UMR 8145, 45, rue des Saints-Pères, 75006, Paris, France

Grégory Nuel, LPSM, CNRS UMR 8001, 4, Place Jussieu, 75005, Paris, France

Olivier Bouaziz, MAP5, CNRS UMR 8145, 45, rue des Saints-Pères, 75006, Paris, France

and use a weighting procedure to combine the three models. Different constraints have also been proposed to make the age-period-cohort model identifiable. However, as noticed by Heuer ([3]), the obtained estimates highly depend on the choice of the constraints. Holford [4] proposed to directly estimate the linear trends of each effect. This procedure leads to results that are difficult to interpret. See Carstensen [5] for a detailed discussion of the identifiability problem of the age-period-cohort model. More recently, Kuang et al. [6] proposed to estimate the second order derivatives of the three effects. This model is implemented in the package `apc` Nielsen [7]. Finally, Carstensen [5] proposed to first fit one submodel (say age-cohort) and then to fit the period effect over the residuals of the first model. This model is implemented in the R package `Epi` [8, 9].

All these approaches can be viewed as parametric models, where the parameters are the age, period, and cohort vector parameters. As such they are also restrictive because they do not allow for interactions between the three effects, that is they assume that one effect does not depend on the other effect's value. A different approach consists in considering the hazard rate as a function of age and either period or cohort and to estimate this bi-dimensional function in a non-parametric setting. No specific structure of the hazard rate is assumed. However, for moderate sample sizes, non-parametric approaches such as the maximum likelihood estimator (MLE) are prone to overparametrization. As a matter of fact the MLE can only be used if a bi-dimensional grid (e.g. of cohort and age intervals) is provided. Without an appropriate method, this grid needs to be arbitrarily chosen. If the number of intervals is too large, the MLE will display a large variance. On the other hand, a too small number of intervals will result in a bias if those intervals are not optimally chosen. Consequently, regularized methods have been proposed in order to avoid overfitting in this non-parametric context. A kernel-type estimator was proposed by Beran [10] and McKeague and Utikal [11] where the cumulative hazard is smoothed using a kernel function. See Keiding [12] for a thorough discussion of methods for hazard inference in age-period-cohort analysis. More recently, Currie and Kirkby [13] proposed a spline estimation procedure to infer the hazard rate as a function of two variables. The authors use a generalized linear model using B-splines and overfitting is dealt with using a penalization over the differences of adjacent splines' coefficients.

In this article, we propose a new non-parametric method for bi-dimensional hazard rate estimation. As the previous non-parametric approaches, this model considers the estimation of the hazard rate with respect to two variables, i.e. either age-cohort, age-period, or period-cohort, without assuming any specific structure on the hazard rate. Inference is made in two dimensions, but through the linear relationship $\text{period} = \text{age} + \text{cohort}$, the hazard rate can be represented as a function of any two of the three variables. Finally, in order to take into account the issue of overfitting, we use a sparsity-inducing penalized likelihood method called adaptive ridge. This iterative method is an approximation of the L_0 norm penalty which makes the computation tractable. We note that the L_0 "norm", defined by $\|\mathbf{u}\|_0 = \#\{j|u_j \neq 0\}$ is not a proper norm but we nevertheless use the term " L_0 norm" hereafter following the notations of e.g. Candès et al. [14]. The method was first introduced by Chartrand and Yin [15] in the context of sparse sensing and applied by Rippe et al. [16] and Frommlet and Nuel [17] in the context of linear regression. It has been used in the context of piecewise constant hazard rate estimation by Bouaziz and Nuel [18]. The present work makes use of this method to perform a segmentation of the hazard rate into constant areas. The novelty of this method lies in the parsimonious representation of the bi-dimensional hazard rate into segmented areas. In particular, the method can efficiently exhibit cohort, age, or period effects, that is, specific changes of the hazard rate due to the date of birth, the age or the calendar time. The penalized likelihood framework used here can also be used to estimate the L_2 norm penalization, which will induce a smoothed estimate of the hazard in a similar way as the aforementioned non-parametric methods.

Our model is introduced in Section 2. The regularization method is then presented in Section 3. In Section 4, the selection of the penalty parameter is discussed. Finally, the performance of our model is assessed through a simulation study in Section 5 and illustrations on the SEER cancer dataset is provided in Section 6.

2 Fused regularized estimation

2.1 Modelization

In the age-period-cohort setting, the date of birth (the cohort) U of each individual is available and the variable of interest is a time-to-event variable of this individual denoted T . The data are subject to right-censoring and are represented as tabulated data over the J cohort intervals and the K age intervals $[c_0, c_1), [c_1, c_2), \dots, [c_{J-1}, c_J)$ and $[d_0, d_1), [d_1, d_2), \dots, [d_{K-1}, d_K)$ respectively, with the convention $c_0 = d_0 = 0$ and $c_J = d_K = \infty$. On a sample of n individuals, the available data can then be rewritten in terms of the exhaustive statistics $\mathbf{O} = (O_{1,1}, \dots, O_{J,K}), \mathbf{R} = (R_{1,1}, \dots, R_{J,K})$, where for $j = 1, \dots, J, k = 1, \dots, K, O_{j,k}$ represents the number of observed events that occurred in the j th cohort interval $[c_{j-1}, c_j)$ and k th age interval $[d_{k-1}, d_k)$ and $R_{j,k}$ represents the total time individuals were at risk in this j th cohort and k th age interval. In the case of register data, the discretization $(c_j), (d_k)$ is imposed by the data and the available data is directly \mathbf{R} and \mathbf{O} , which are often called the *cases* and *person-years*, respectively. See for instance Carstensen [5] for an example of such data. The aim is to estimate the hazard rate, defined as:

$$\lambda(t|u) = \lim_{dt \rightarrow 0} \frac{1}{dt} P(t \leq T < t + dt | T \geq t, U = u).$$

In the age-cohort setting $\lambda(t|u)$ is assumed to be piecewise constant:

$$\lambda(t|u) = \sum_{j=1}^J \sum_{k=1}^K \lambda_{j,k} \mathbb{1}_{[c_{j-1}, c_j) \times [d_{k-1}, d_k)}(t, u),$$

and inference is made over the $J \times K$ dimension parameter $\lambda = (\lambda_{1,1}, \dots, \lambda_{J,K})$. Note that the hazard can be equivalently defined as a function of age and period or as a function of period and cohort where the period is defined as the calendar time, that is: period = cohort + age. For illustration, the change of coordinates between the age-period and age-cohort diagrams is represented in Figure 1. In our models, the hazard will be considered as a function of solely age and cohort since the influence of any of the two elements of age, period or cohort can be retrieved using this reparametrization.

2.2 Penalized likelihood

Following Aalen et al. ([19]) the negative log-likelihood takes the form

$$\ell_n(\lambda) = \sum_{j=1}^J \sum_{k=1}^K \{ \lambda_{j,k} R_{j,k} - O_{j,k} \log(\lambda_{j,k}) \}. \quad (1)$$

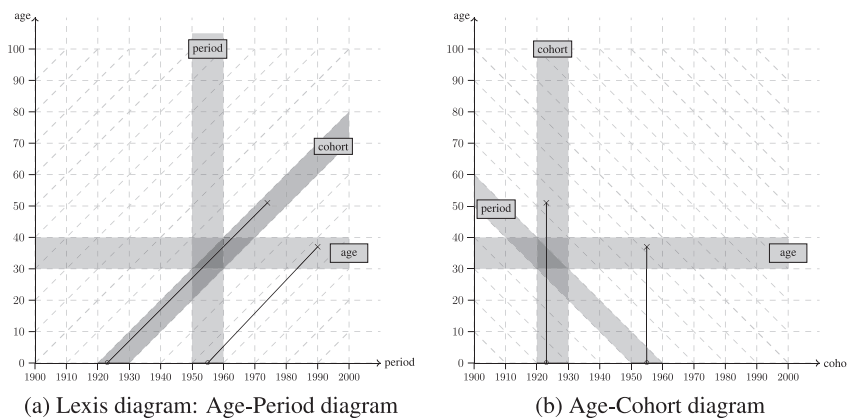


Figure 1: Diagrams representing the lives of individuals: in the age-period plane (a) – called Lexis diagram – and in the age-cohort plane (b). Solid lines represent lives of individuals until occurrence of the event of interest. The same age, cohort, and period intervals are displayed in gray.

The authors also noticed that this log-likelihood is equivalent to a log-likelihood arising from a Poisson model. However, note that no distribution assumptions are made on the data and in particular the $O_{j,k}$ are not assumed to be Poisson distributed (see Carstensen [5]), for a discussion on the ‘‘Poisson’’ model). Minimizing ℓ_n yields an explicit maximum likelihood estimate $\hat{\lambda}_{j,k}^{\text{mle}} = O_{j,k}/R_{j,k}$. However, for moderate sample sizes this estimator is overfitted, especially in places of the age-cohort plane where few events are recorded. To remedy this problem we propose in the following to penalize the differences between adjacent values of the hazard in the log-likelihood.

For computation convenience, we first reparametrize the model: $\eta_{j,k} = \log \lambda_{j,k}$, for $1 \leq j \leq J$ and $1 \leq k \leq K$. The goal of this work is to estimate the minimizer of the function

$$\ell_n(\boldsymbol{\eta}) + \frac{\kappa}{2} \sum_{j=1}^{J-1} \sum_{k=1}^K \|\eta_{j+1,k} - \eta_{j,k}\|_0 + \frac{\kappa}{2} \sum_{j=1}^J \sum_{k=1}^{K-1} \|\eta_{j,k+1} - \eta_{j,k}\|_0, \quad (2)$$

where $\ell_n(\boldsymbol{\eta})$ was defined in Eq. (1) and κ is a penalty constant used as a tuning parameter. In the previous equation, the L_0 norm penalty over the differences of adjacent parameter values yields a piecewise constant estimate of the hazard rate. However this function is not tractable to minimize due to the L_0 norm. Hence we use the adaptive ridge, which, as is explained in Section 3, can be seen as an approximate solution to this problem.

Let us define the weighted L_2 penalized model:

$$\ell_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w}) = \ell_n(\boldsymbol{\eta}) + \frac{\kappa}{2} \sum_{j=1}^{J-1} \sum_{k=1}^K v_{j,k} (\eta_{j+1,k} - \eta_{j,k})^2 + \frac{\kappa}{2} \sum_{j=1}^J \sum_{k=1}^{K-1} w_{j,k} (\eta_{j,k+1} - \eta_{j,k})^2, \quad (3)$$

where $\boldsymbol{v} = (v_{1,1}, \dots, v_{J-1,K})$ and $\boldsymbol{w} = (w_{1,1}, \dots, w_{J,K-1})$ are constant positive weights of respective dimensions $(J-1)K$ and $J(K-1)$. Note that the case $\kappa = 0$ corresponds to the maximum likelihood estimation and the case $\kappa = \infty$ corresponds to a hazard uniformly constant over the age and cohort intervals. The parameter κ needs to be chosen in an appropriate way in order to obtain a compromise between these two extreme situations. This is addressed in Section 4.

This model does not attempt to estimate the age, period and cohort effect as parameter vectors. Instead, it performs a regularized estimation of λ that has no age-period-cohort-type structure.

In the next section, we introduce an algorithm to minimizing Eq. (3), which will be used for estimating both L_2 and L_0 penalties. We then introduce the estimation procedures for both fused L_2 and L_0 penalties.

3 Numerical optimization

In this section, we first introduce the weighted L_2 penalized negative log-likelihood and derive how to minimize it. Then, two different expressions of the weights \boldsymbol{v} and \boldsymbol{w} are proposed which correspond to two different types of regularization of the hazard rate. The first one implements the adaptive ridge and yields a piecewise constant estimate. The second one uses constant weights and yields a smooth estimate.

3.1 Fused L_2 penalty estimate

Minimization of ℓ_n^κ is performed using the Newton–Raphson method (see Algorithm 1). Let $\mathbf{U}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w}) = \partial \ell_n^\kappa / \partial \boldsymbol{\eta}$ be the gradient of the penalized negative log-likelihood and $\mathbf{I}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w}) = \partial \mathbf{U}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w}) / \partial \boldsymbol{\eta}^T$ be its Hessian matrix.

For $1 \leq j, j' \leq J$ and $1 \leq k, k' \leq K$, simple algebra yields

$$\frac{\partial \ell_n^\kappa(\boldsymbol{\eta})}{\partial \eta_{j,k}} = \exp(\eta_{j,k}) R_{j,k} - O_{j,k}, \quad \frac{\partial^2 \ell_n^\kappa(\boldsymbol{\eta})}{\partial \eta_{j',k'} \partial \eta_{j,k}} = \mathbb{1}_{j=j', k=k'} \exp(\eta_{j,k}) R_{j,k}, \quad \text{and}$$

$$\begin{aligned} \frac{\partial \ell_n^\kappa(\boldsymbol{\eta})}{\partial \eta_{j,k}} &= \frac{\partial \ell_n(\boldsymbol{\eta})}{\partial \eta_{j,k}} + \kappa \left[-v_{j,k} (\eta_{j+1,k} - \eta_{j,k}) + v_{j-1,k} (\eta_{j,k} - \eta_{j-1,k}) \right] \\ &\quad + \kappa \left[-w_{j,k} (\eta_{j,k+1} - \eta_{j,k}) + w_{j,k-1} (\eta_{j,k} - \eta_{j,k-1}) \right], \\ \frac{\partial^2 \ell_n^\kappa(\boldsymbol{\eta})}{\partial \eta_{j',k'} \partial \eta_{j,k}} &= \frac{\partial^2 \ell_n(\boldsymbol{\eta})}{\partial \eta_{j',k'} \partial \eta_{j,k}} + \kappa \left[\mathbb{1}_{j=j',k=k'} (v_{j',k'} + v_{j'-1,k'} + w_{j',k'} + w_{j',k'-1}) \right. \\ &\quad \left. - v_{j',k'} \mathbb{1}_{j=j'+1,k=k'} - v_{j'-1,k'} \mathbb{1}_{j=j'-1,k=k'} - w_{j',k'} \mathbb{1}_{j=j',k=k'+1} - w_{j',k'-1} \mathbb{1}_{j=j',k=k'-1} \right]. \end{aligned}$$

From the last equation, the Hessian matrix can be written

$$\mathbf{I}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w}) = \frac{\partial^2 \ell_n(\boldsymbol{\eta})}{\partial \boldsymbol{\eta} \partial \boldsymbol{\eta}^T} + \kappa \mathbf{B}(\boldsymbol{\eta}),$$

where $\mathbf{B}(\boldsymbol{\eta})$ is a band matrix of bandwidth equal to $\min(J, K) - 1$. Thus the Hessian matrix is also a band matrix of bandwidth $\min(J, K) - 1$. Using Cholesky decomposition, the computation of $\mathbf{I}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w})^{-1} \mathbf{U}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w})$ has a $\mathcal{O}(\min(J, K)JK)$ complexity instead of $\mathcal{O}(J^3K^3)$.

Algorithm 1: Newton–Raphson procedure with constant weights

```

1: function NEWTON_RAPHSON( $\mathbf{O}, \mathbf{R}, \kappa, v, w$ )
2:    $\boldsymbol{\eta} \leftarrow \mathbf{0}$ 
3:   while not converge do
4:      $\boldsymbol{\eta}^{\text{new}} \leftarrow \boldsymbol{\eta} - \mathbf{I}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w})^{-1} \mathbf{U}_n^\kappa(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w})$ 
5:      $\boldsymbol{\eta} \leftarrow \boldsymbol{\eta}^{\text{new}}$ 
6:   end while
7:   return  $\boldsymbol{\eta}$ 
8: end function

```

A ridge-type penalization is performed when setting $\boldsymbol{v} = \boldsymbol{w} = \mathbf{1}$ in Eq. (2). In this case the penalization corresponds to the square of the first-order differences of $\boldsymbol{\eta}$, which yields a smooth estimator of the hazard rate. This estimate is obtained directly from Algorithm 1.

We make a note that Eq. (3) allows for some flexibility in the regularization. Indeed, one could set different values to \boldsymbol{v} and \boldsymbol{w} to manually tune the importance of the regularization between different regions of the plane and between the two variables.

Finally, note that this method yields an estimate similar to the spline method of Ogata and Katsura [20], who penalizes over the second-order differences instead of the first-order differences. This means that for arbitrarily large values of the penalty constant, the regularized hazard will be a constant function instead of a linear function.

3.2 Fused adaptive ridge estimate

In this section, we derive a computationally tractable procedure to minimize Eq. (2).

We make use of the adaptive ridge, which minimizes a non-convex penalty by iteratively minimizing approximations of the penalty. The adaptive ridge can be used to approximate any L_q penalty ($0 < q < 1$) and it extends to the case $q = 0$, the latter case corresponding to the logarithmic penalty in lieu of the L_0 penalty. More precisely, the L_0 adaptive ridge makes use of the fact that when ε is close to zero, the function

$$\ell_n(\boldsymbol{\eta}) + \frac{\kappa}{2} \sum_{j=1}^{J-1} \sum_{k=1}^K \log \left(1 + \frac{(\eta_{j+1,k} - \eta_{j,k})^2}{\varepsilon^2} \right) + \frac{\kappa}{2} \sum_{j=1}^J \sum_{k=1}^{K-1} \log \left(1 + \frac{(\eta_{j,k+1} - \eta_{j,k})^2}{\varepsilon^2} \right)$$

is a good approximation of (2) (up to a multiplicative constant on the penalty terms). The adaptive ridge consists in a numerical scheme minimizing the latter function. It does so by iteratively solving L_2 penalty problems (hence its name), and is thus simple to implement. Details are given in Section S1 of the Supplementary Materials.

As pointed out by a reviewer, another iterative penalized method [21] could have been used which iteratively solves L_1 penalty problems to approximate the L_0 penalty. As explained by Wipf and Nagarajan [22], these two methods are very similar in that they both minimize a logarithm penalized problem using two different approximations. In Section S1 of the Supplementary Materials, we make explicit the link between the two methods and show that they belong to the same class of optimization schemes. We refer to works related to either of the two methods in Section S1.3.

We implement the adaptive ridge procedure by minimizing Eq. (3) with the weights adapted iteratively. We iterate between updating

$$\boldsymbol{\eta}^{\text{new}} \leftarrow \arg \min_{\boldsymbol{\eta}} \ell_n^{\kappa}(\boldsymbol{\eta}, \boldsymbol{v}, \boldsymbol{w})$$

using Algorithm 1 and updating the values of the weights:

$$\begin{cases} v_{j,k}^{\text{new}} = \left(\left(\eta_{j+1,k}^{\text{new}} - \eta_{j,k}^{\text{new}} \right)^2 + \varepsilon_v^2 \right)^{-1}, \\ w_{j,k}^{\text{new}} = \left(\left(\eta_{j,k}^{\text{new}} - \eta_{j,k-1}^{\text{new}} \right)^2 + \varepsilon_w^2 \right)^{-1}, \end{cases}$$

where ε_v and ε_w are constants negligible compared to 1.

We now elaborate on the estimation procedure. The algorithm is said to converge if the absolute difference of all weighted differences in $\eta_{j,k}$ are below a given threshold (we use 10^{-8} in our implementation). At convergence, $v_{j,k} (\eta_{j+1,k} - \eta_{j,k})^2$ will be either very close to 0 if $|\eta_{j+1,k} - \eta_{j,k}|$ is smaller than ε_v or very close to 1 if $|\eta_{j+1,k} - \eta_{j,k}|$ is greater than ε_v — and similarly for $w_{j,k} (\eta_{j,k+1} - \eta_{j,k})^2$. We then set them to 0 or 1 using a thresholding, so that values smaller than 0.99 are set to 0 and values larger than 0.99 are set to 1 (in practice, the value of this threshold has little effect, since at convergence the weighted differences are distant to 0 or 1 by $\sim 10^{-7}$).

As with other penalized methods and as pointed out in Frommlet and Nuel [17], the adaptive ridge penalization scheme induces a shrinkage bias. Therefore, after segmentation of the $\eta_{j,k}$ s, the hazard rate is estimated on each constant area using the unpenalized maximum likelihood estimator. These constant areas are defined as connected components of a graph. We first create the graph whose vertices are the JK age-cohort rectangles and whose edges are the connections between adjacent cells that have differences equal to 0. Then, each connected component of this graph is a different area over which the hazard has been estimated to be constant. The extraction of connected components from the graph is done using the package `igraph` [23]. The log-hazard $\eta^{(r)}$ of the r th constant area is such that $\forall [c_{j-1}, c_j] \times [d_{k-1}, d_k] \in r, \eta_{j,k} = \eta^{(r)}$. The values of $\eta^{(r)}$ are then estimated in a second step, using unpenalized maximum likelihood estimation: $\hat{\eta}^{(r)} = \log(O^{(r)}/R^{(r)})$ where $O^{(r)}$ is the number of events in the r th constant area and $R^{(r)}$ is the time at risk in the r th constant area.

This algorithmic procedure is summarized in Algorithm 2. In practice, the stopping criterion for the adaptive ridge algorithm is when the absolute difference between successive values of the weighted differences is smaller than a predefined value — we use 10^{-8} in our implementation. Moreover, following Frommlet and Nuel [17], we have set $\varepsilon_v = \varepsilon_w = 10^{-5}$.

Algorithm 2: Adaptive ridge procedure

```

1: function ADAPTIVE RIDGE( $\mathbf{O}, \mathbf{R}, \kappa$ )
2:    $\boldsymbol{\eta} \leftarrow \mathbf{0}$ 
3:    $\mathbf{v} \leftarrow \mathbf{1}$ 
4:    $\mathbf{w} \leftarrow \mathbf{1}$ 
5:   while not converge do
6:      $\boldsymbol{\eta}^{\text{new}} \leftarrow \text{NEWTON RAPHSON}(\mathbf{O}, \mathbf{R}, \kappa, \mathbf{v}, \mathbf{w})$ 
7:      $v_{j,k}^{\text{new}} \leftarrow \left( \left( \eta_{j+1,k}^{\text{new}} - \eta_{j,k}^{\text{new}} \right)^2 + \varepsilon_v^2 \right)^{-1}$ 
8:      $w_{j,k}^{\text{new}} \leftarrow \left( \left( \eta_{j,k}^{\text{new}} - \eta_{j,k-1}^{\text{new}} \right)^2 + \varepsilon_w^2 \right)^{-1}$ 
9:      $\boldsymbol{\eta} \leftarrow \boldsymbol{\eta}^{\text{new}}$ 
10:  end while
11:  Compute ( $\mathbf{O}^{\text{new}}, \mathbf{R}^{\text{new}}$ ) selected ( $\boldsymbol{\eta}, \mathbf{v}^{\text{new}}, \mathbf{w}^{\text{new}}$ )
12:   $\boldsymbol{\eta}^{\text{new}} \leftarrow \log(\mathbf{O}^{\text{new}} / \mathbf{R}^{\text{new}})$ 
13:  return  $\boldsymbol{\eta}^{\text{new}}$ 
14: end function

```

4 Choice of the penalty parameter κ

In practice, the hazard rate needs to be estimated for a set of penalty constants and the choice of κ is determined as the penalty that provides the best compromise between model fit and reduced variability of the hazard rate estimate. For the L_0 regularization model, different values of the penalty constant lead to different segmentations of the $\eta_{j,k}$. As a consequence, the problem of choosing the optimal penalty constant can be rephrased as the problem of choosing the optimal model among a set of models $\mathcal{M}_1, \dots, \mathcal{M}_M$, where each of these models corresponds to a different segmentation of the $\eta_{j,k}$ and M is the maximum number of different models. In this section we propose different methods to select the optimal model. Comparison of the efficiency of the different methods will be analyzed in Section 5 on simulated data.

We recall that \mathbf{R} and \mathbf{O} are the exhaustive statistics and $\boldsymbol{\eta}$ is the parameter to be estimated in our two models. Bayesian criteria attempt to maximize the posterior probability $P(\mathcal{M}_m | \mathbf{R}, \mathbf{O}) \propto P(\mathbf{R}, \mathbf{O} | \mathcal{M}_m) \pi(\mathcal{M}_m)$, where $P(\mathbf{R}, \mathbf{O} | \mathcal{M}_m)$ is the integrated likelihood and $\pi(\mathcal{M}_m)$ is the prior distribution on the model. This problem is equivalent to minimizing $-2 \log P(\mathcal{M}_m | \mathbf{R}, \mathbf{O})$. By integration

$$P(\mathbf{R}, \mathbf{O} | \mathcal{M}_m) = \int_{\boldsymbol{\eta}} P(\mathbf{R}, \mathbf{O} | \mathcal{M}_m, \boldsymbol{\eta}) \pi(\boldsymbol{\eta}) d\boldsymbol{\eta},$$

where $P(\mathbf{R}, \mathbf{O} | \mathcal{M}_m, \boldsymbol{\eta})$ is the likelihood and $\pi(\boldsymbol{\eta})$ is the prior distribution of the parameter, which is taken constant in the following. Thus Bayesian criteria are defined as

$$-2 \log(P(\mathcal{M}_m | \mathbf{R}, \mathbf{O})) = 2\ell_n(\hat{\boldsymbol{\eta}}_m) + q_m \log n - 2 \log \pi(\mathcal{M}_m) + \mathcal{O}_P(1),$$

where q_m is the dimension of the model \mathcal{M}_m i.e., the number of constant areas selected by the adaptive ridge algorithm.

The BIC [24] corresponds to the Bayesian criterion obtained when one neglects the term $\pi(\mathcal{M}_m)$, which is equivalent to having a uniform prior on the model:

$$\text{BIC}(m) = 2\ell_n(\hat{\boldsymbol{\eta}}_m) + q_m \log n.$$

As explained by Żak-Szatkowska and Bogdan [25], a uniform prior on the model is equivalent to a binomial prior on the model dimension $\mathcal{B}(JK, 1/2)$. When the true model's dimension is much smaller than the

maximum possible dimension JK , the BIC tends to give too much importance to models of dimensions around $JK/2$, which will result in underpenalized estimators. To this effect, Chen and Chen [26] have developed an extended Bayesian information criterion called EBIC_0 (or EBIC for short). One can write $\pi(\mathcal{M}_m) = \text{P}(\mathcal{M}_m | \mathcal{M}_m \in \mathcal{M}_{[q_m]}) \text{P}(\mathcal{M}_m \in \mathcal{M}_{[q_m]})$ where $\mathcal{M}_{[q_m]}$ is the set of models of dimension q_m . The EBIC_0 criterion is defined by setting $\text{P}(\mathcal{M}_m | \mathcal{M}_m \in \mathcal{M}_{[q_m]}) = 1 / \binom{JK}{q_m}$ and $\text{P}(\mathcal{M}_m \in \mathcal{M}_{[q_m]}) = 1$. Thus

$$\pi(\mathcal{M}_m) = \binom{JK}{q_m}$$

and

$$\text{EBIC}_0(m) = 2\ell_n(\hat{\eta}_m) + q_m \log n + 2 \log \binom{JK}{q_m}.$$

Note that the EBIC_0 assigns the same *a priori* probability to all models of same dimension. Therefore, when the true model's dimension is not close to $JK/2$ the EBIC_0 will be able to select this model more easily. Namely, when the true model's dimension is very small the EBIC_0 will tend to choose very sparse models.

The last criterion that will be used is the Akaike information criterion [27], or AIC, defined as $\text{AIC}(m) = 2\ell_n(\hat{\eta}_m) + 2q_m$. This criterion is known for performing better than the BIC in terms of mean squared error, however the BIC will tend to select sparser models than the AIC.

Note that Bayesian criteria and the AIC can only be used for the L_0 regularized estimation only, since the L_2 model does not perform a model selection. An alternative to performing model selection is to use the K-fold cross validation. With this method, the data are split at random into L parts. The estimated parameter obtained when the l th part is left out is noted $\hat{\eta}^{-l}(\kappa)$ and the cross-validated score is defined as

$$\text{CV}(\kappa) = \sum_{l=1}^L \ell_n^{\kappa,l}(\hat{\eta}^{-l}),$$

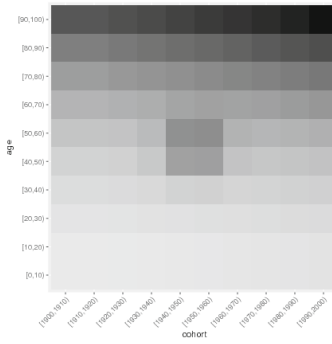
where $\ell_n^{\kappa,l}$ is the negative log-likelihood evaluated on the l th part of the data. The optimal penalty constant is obtained by minimizing $\text{CV}(\kappa)$ with respect to κ . The L -fold cross validation method can be used for both the L_0 regularized estimation and the L_2 regularized estimation. However, this method is numerically time consuming as the estimator has to be computed L times while Bayesian criteria or the AIC provide direct methods to perform model selection from the original estimator. In the simulation studies and data analysis, we set $L = 10$.

5 Simulation study


5.1 Simulation designs

In this section, our piecewise estimation method is compared with the AGE-COHORT model and with the L_2 penalty estimate. The different criteria for model selection are also compared with each other. We present two simulation designs. In the first one, the true hazard rate is generated from a smooth age-cohort model which includes an interaction term on a small region of the age-cohort plane. In the second case, the true hazard rate is a piecewise constant function with four heterogeneous areas. The two true hazards are displayed in Figure 2, both in greyscale and in perspective plot.


The simulation design is as follows. We set $J = 10$ equally spaced age intervals ($[0, 10)$, ..., $[90, 100]$) and $K = 10$ equally spaced cohort intervals ($[1900, 1910)$, ..., $[1990, 2000]$). In order to simulate a dataset, the cohorts are first sampled uniformly over the 10 cohort intervals and the age is then simulated using the corresponding hazard. Censoring is then simulated as a uniform distribution over the age interval $[75, 100]$ for all cohorts such that all observed events are comprised in the age interval $[0, 100]$. Since in practice one does not know the appropriate discretization in advance, a different discretization was used for the estimation procedure: the age and cohort intervals were defined as 5-year length intervals instead of 10 for the true




(a) Smooth true hazard – heatmap



(b) Smooth true hazard – perspective



(c) Piecewise constant true hazard – heatmap



(d) Piecewise constant true hazard – perspective

Figure 2: True hazard of the two simulation designs: smooth hazard in heatmap (a) and perspective plot (b) and piecewise constant hazard in heatmap (c) and perspective plot (d).

hazard. As a result, a total of 20×20 parameters need to be estimated. We simulated data of sample sizes 100, 400, 1000, 4000, and 10,000. For each sample size, the simulation and estimation were replicated 500 times.

5.1.1 Smooth true hazard

The smooth true hazard (Figure 2(a) and (b)) is generated using the age-cohort model $\log \lambda_{j,k} = \mu + \alpha_j + \beta_k$ with an intercept $\mu = \log(10^{-2})$. The age effect vector α and cohort effect vector β are arithmetic sequences such that $\alpha_1 = 0$, $\alpha_7 = 2.5$, $\beta_1 = 0$, and $\beta_K = 0.3$. An interaction term is added to the hazard. It corresponds to a bump in the hazard located in the neighbourhood of the region of the age-cohort plane (45, 1945). The bump is defined as 10 times the Gaussian density function with mean (1945, 45) and with a diagonal variance-covariance matrix with diagonal equal to (50, 50). This true hazard displays a sharp increase for high values of the age, which implies that few events will be recorded in this region. On average, 91% of the events are observed in this simulation design.

5.1.2 Piecewise constant true hazard

The piecewise constant true hazard (Figure 2(c) and (d)) has four constant areas over the age-cohort square $[0, 100] \times [1900, 2000]$. On average, 71% of the events are observed in this simulation design.

5.2 Performance of the estimation methods in terms of MSE

Our two estimation methods (L_0 penalty and L_2 penalty) are compared in terms of the mean squared error (MSE) in each simulation scenario. The different selection methods for the penalty (AIC, BIC, EBIC and cross-validation) are included. We compare our methods with the maximum likelihood estimate (MLE), which serves as baseline for comparison. The results are presented in Table 1, which reports the relative mean square errors with respect to the MLE for easier comparison.

Table 1: Relative mean squared errors of the L_0 and L_2 methods with respect to the maximum likelihood estimate (MLE), for different sample sizes and different estimation methods.

Sample size	L_0 method				L_2 method	MLE
	AIC	BIC	EBIC	CV	CV	
(a) Smooth true hazard						
100	1.016	0.988	0.011	0.011	0.002	1
400	1.005	0.845	0.144	0.026	0.004	1
1000	0.946	0.628	0.024	0.020	0.006	1
40000	0.851	0.267	0.054	0.037	0.011	1
10000	0.634	0.144	0.113	0.057	0.024	1
(b) Piecewise constant true hazard						
100	1.004	1.001	0.003	0.003	0.002	1
400	0.984	0.775	0.036	0.029	0.012	1
1000	0.829	0.408	0.092	0.085	0.024	1
4000	0.715	0.128	0.090	0.110	0.058	1
10000	0.720	0.083	0.065	0.081	0.107	1

For easier comparison, the mean squared errors are given as the ratio with respect to the mean squared error of the MLE. Panel (a): smooth true hazard. Panel (b): piecewise constant true hazard.

Overall, the EBIC and cross-validated criteria outperform the AIC and the BIC for the two simulations scenarios. This is particularly true for small sizes where the AIC and the BIC behave very poorly. As expected, the L_2 penalty estimator is the most performant of all estimators in the smooth true hazard scenario (Table 1(a)) and the L_0 method performs better in the piecewise constant hazard scenario (Table 1(b)) than in the smooth true hazard scenario. The L_2 norm estimator is also the most performant of all estimators in the piecewise constant hazard scenario except for very large sample sizes ($n = 10,000$) where the BIC, EBIC and cross-validated criterion provide slightly better performances. In both scenarios, the EBIC always outperforms the AIC, the BIC and the cross-validated criterion.

Finally, note that both the L_2 penalty and the L_0 penalty (with the EBIC or cross-validation) vastly outperform the MLE. This validates that our penalized approach is useful to reduce the mean square error of the estimate. The degree to which the penalized methods outperform the MLE decreases as the sample size increases, but they still outperform the MLE for a sample size of 10,000.

Different censoring rates were also studied which showed a degradation of the performances of the overall estimators as the percentage of censored events increases. The performance in terms of number of selected areas was also investigated. It showed that the EBIC and CV criterion perform better at selecting sparse models with few areas, while the AIC and BIC tend to overestimate the true number of areas. Indeed, for sample size 4000, the 80% inter-quantile range of the selected number of areas is [3, 5] for the EBIC and [1, 5] for the CV, whereas it is [3, 13] and [36, 72] for the BIC and AIC respectively. These experiments are not reported here.

In conclusion, the simulation experiments suggest to use the EBIC among all different criteria for the L_0 penalty as it provides the best tradeoff between computation time and estimation performance. It has been shown that using the L_0 penalty is beneficial even when the true hazard is not piecewise constant, as our simulations show that the performance of this estimate exceeds that of the MLE with a smooth true hazard.

5.3 Perspective plots of the estimation methods

In this section the performance of our two estimates is assessed visually by comparison with the true hazard. The standard age-cohort model [4] has also been implemented. This model assumes that the hazard has the

following expression:

$$\log \lambda_{j,k} = \mu + \alpha_j + \beta_k,$$

where μ is the intercept, α is the age effect and β is the cohort effect. It should be noted that this model does not allow for interactions between age and cohort effects. Perspective plots of the median hazard estimations over 500 replications are presented in Figures 3 and 4 for the smooth and piecewise constant true hazard respectively. For the L_0 regularized estimate, the penalty constant is chosen using the EBIC.

In Figure 3, it is seen that the age-cohort model is not able to estimate the central bump in the hazard. On the contrary, the smoothed estimate accurately recovers the shape of the true hazard except for the high

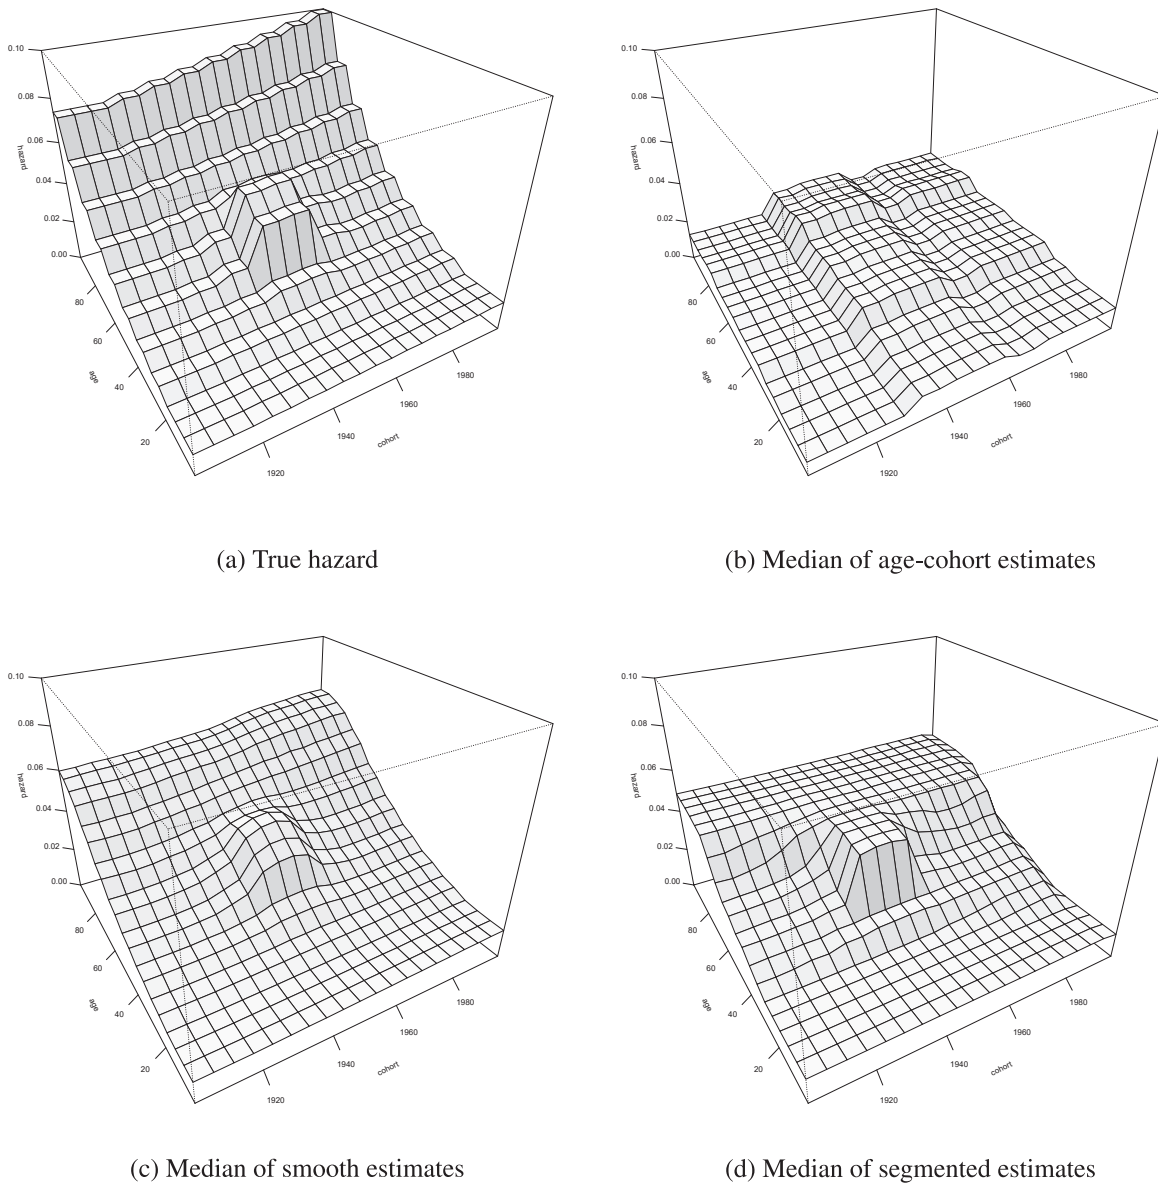


Figure 3: Smooth true hazard and corresponding estimates. The sample size is 4000 and the hazard estimates are medians taken over 500 simulations. The estimations are performed in the age-cohort plane and with different methods. Panel (a) represents the true hazard used to generate the data, panel (b) represents the hazard estimated using the age-cohort model, panel (c) represents the smoothed estimate, and panel (d) represents the segmented estimate with the EBIC criterion.

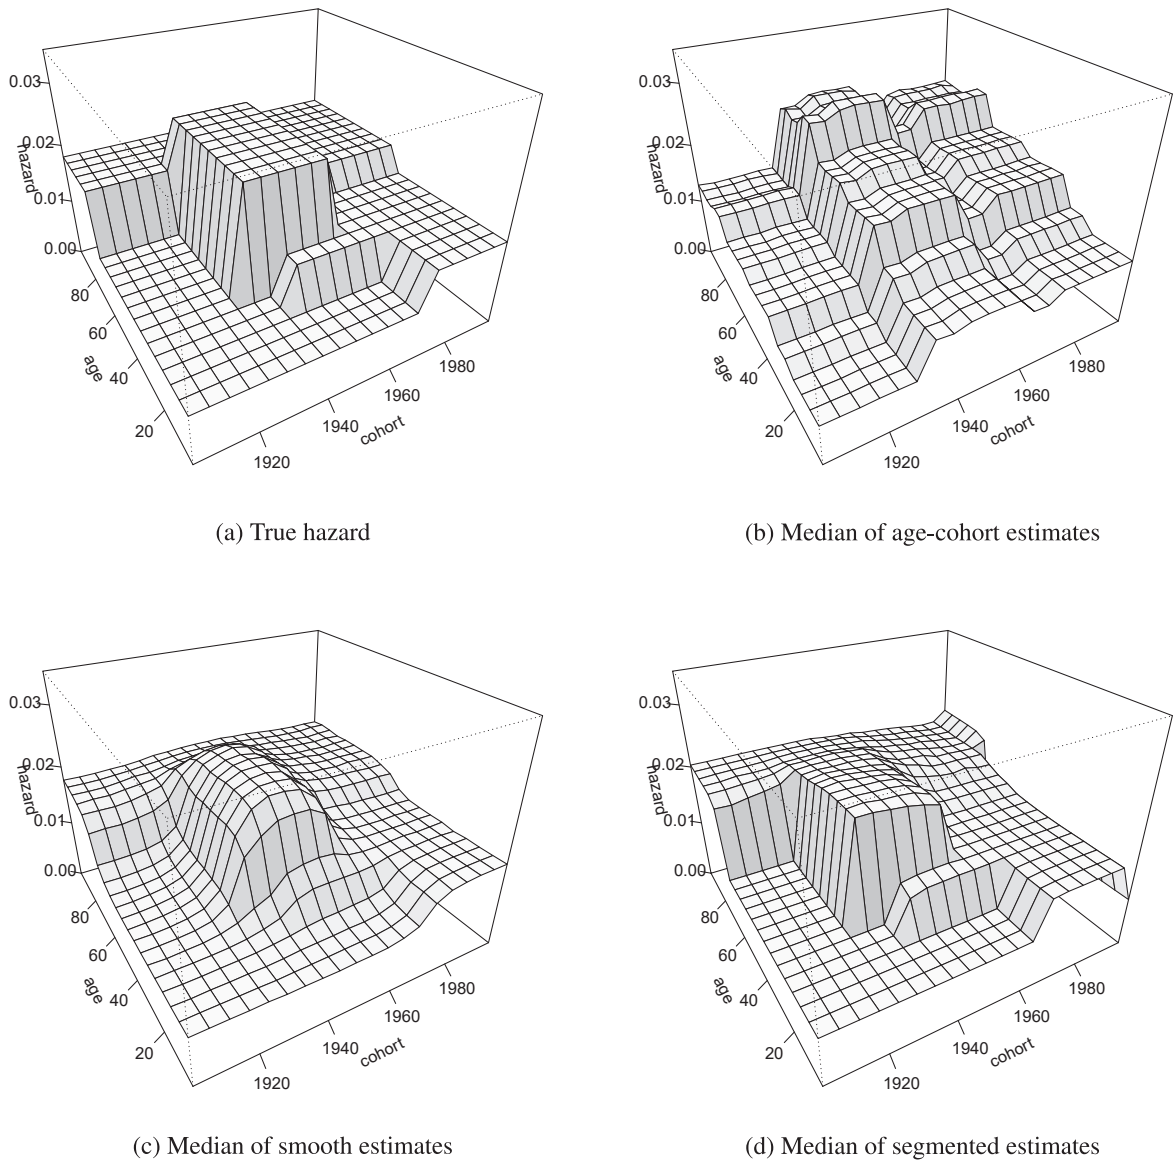


Figure 4: Piecewise constant true hazard and corresponding estimates. The sample size is 4000 and the hazard estimates are medians taken over 500 simulations. The estimations are performed in the age-cohort plane and with different methods. Panel (a) represents the true hazard used to generate the data, panel (b) represents the hazard estimated using the age-cohort model, panel (c) represents the smoothed estimate, and panel (d) represents the segmented estimate with the EBIC criterion.

values of age where few events are observed. Interestingly, one sees that our segmentation method provides results similar to the smoothing technique even though the true hazard is not piecewise constant.

The results in Figure 5 yield similar conclusions. The age-cohort model behaves very poorly due to its constrained structure while the ridge and adaptive estimates provide satisfactory results. In particular the overall shape of the true hazard is correctly estimated by the L_0 penalty.

6 Real data application

Our method is applied to data of survival times after diagnosis of breast cancer. The dataset is provided by the surveillance, epidemiology, and end results (SEER) program from the US national cancer institute (NCI). SEER collects medical data of cancers (including stage of cancer at diagnosis and the type of tumor) and follow-up

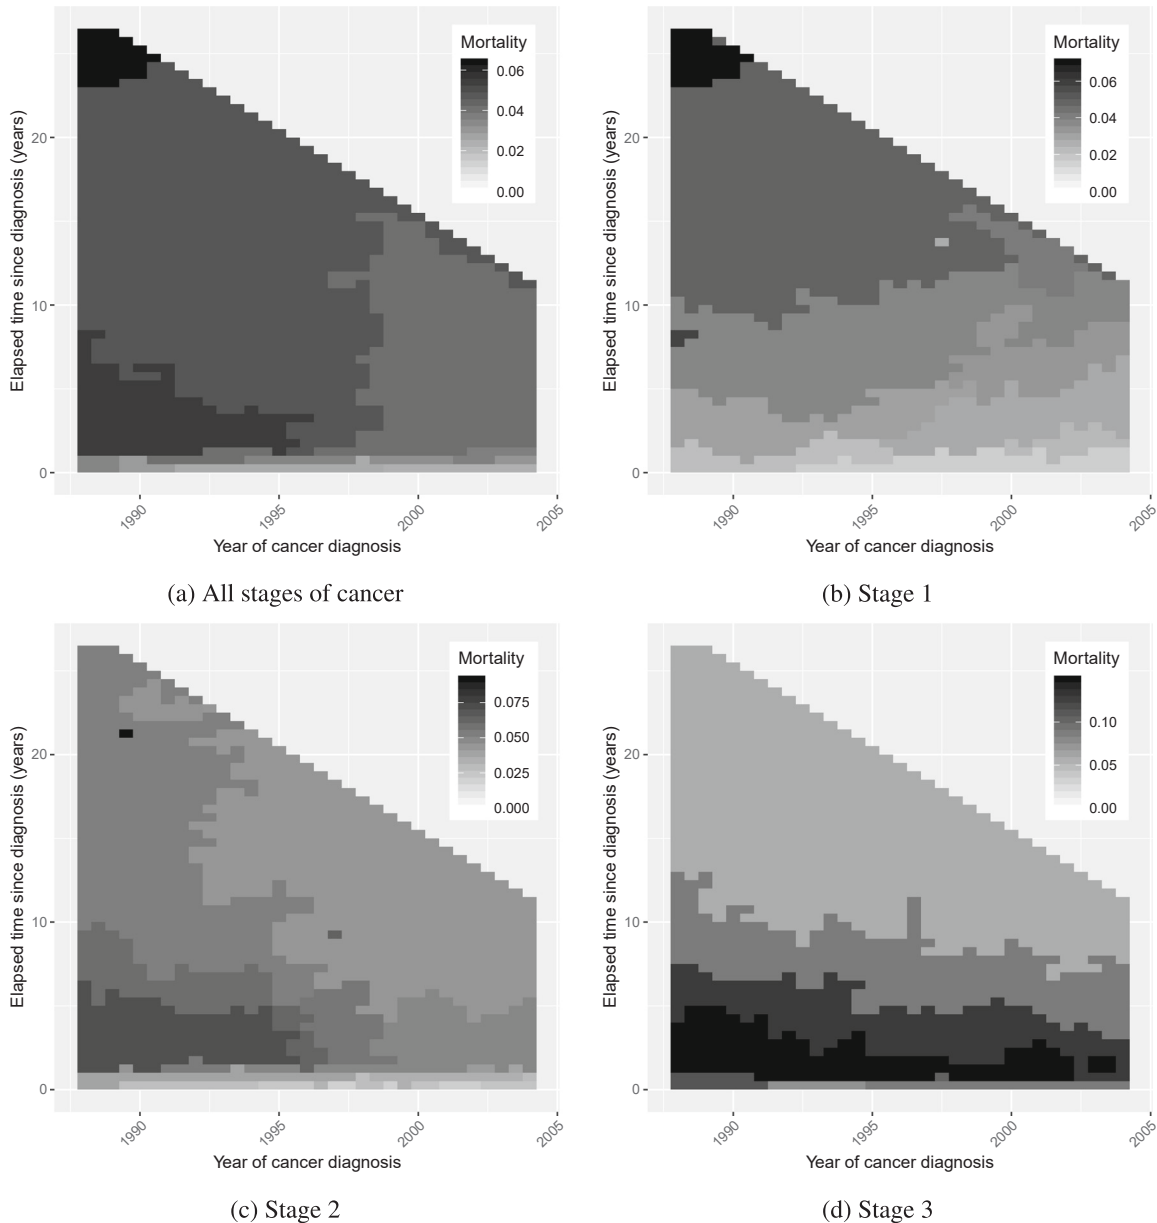


Figure 5: Estimated hazard of death after diagnosis of breast cancer for different stages of cancer. The estimate is obtained with the L_0 regularization. The upper right corner of every graph corresponds to the region where no data are available. Note that the grey-color scales are different between panels.

data of patients in the form of a registry. Around 28 percent of the US population is covered by the program. The registry started in February 1973 and the available current dataset includes follow-up data until January 2015. We refer to the website <https://seer.cancer.gov/> for information about the SEER Program and its publicly available cancer data.

In this study the duration of interest T is the time from breast cancer diagnosis to death in years, the variable U is the date of diagnosis (in years) and the period is the calendar time (in years). Patients continuously entered the study between 1973 and 2015 and right-censoring occurred for patients that were still alive at the end of follow-up or for those that were lost to follow-up.

The breast cancer data was extracted using the package SEERaBomb. For the sake of comparison, the subsample of malignant, non-bilateral breast tumor cancers was extracted from the dataset, such that the

data comprises 1,265,277 women with 60 percent of censored individuals. Times from diagnosis to last day of follow-up vary between 0 and 41 years, and the dates of cancer diagnosis U_i vary between 1973 and 2015. Death from another cause than cancer is available in the dataset and is accounted for as right-censoring.

The implementation of our adaptive ridge method aims at two goals. Firstly we aim at simultaneously detecting a cohort effect and an age effect, that is the evolution of the mortality with respect to the time elapsed since cancer diagnosis (age effect) and with respect to the date of diagnosis (cohort effect). Secondly, our method will provide estimation of the hazard rates on the resulting heterogeneous areas. The method is first applied on the whole sample of 1,265,277 individuals. In order to take into account the fact that mortality from cancer highly depends on the cancer stage, we also perform a stratified analysis with respect to the stage of cancer at diagnosis. For this purpose, we use the cancer stage classification provided by the SEER data: we keep the patients with cancer stages 1, 2, and 3 at the time of diagnosis. This classification closely follows that of the American joint committee on Cancer (AJCC), 3rd edition; the details are given at page 86 of the manual entitled comparative staging guide for cancer, available at <https://seer.cancer.gov>. The main difference between the two classifications is that the SEER Program classifies the cases where lymph node status cannot be assessed as if there was no regional lymph node metastasis.

The L_0 estimates for the whole sample and for each cancer stage are displayed in Figure 5. We see that the different stages of cancer at diagnosis have a great impact on the survival times. For stage 1 cancers, the mortality is low between 0 and 4–5 years after diagnosis, and steadily increases afterwards. The date of diagnosis seems to have no impact on the mortality of stage 1 cancers. On the other hand, stage 2 cancers exhibit a strong effect of the date of diagnosis: around 1995–1997, the mortality significantly decreases. This can correspond to an improvement of the treatment of breast cancer around that period in the United States. Finally, stage 3 cancers display a very high hazard rate across all dates of diagnosis. This seems to indicate that the evolution in treatments of breast cancer had a significant impact on the survival times after diagnosis, but almost exclusively when cancers were diagnosed at stage 2. Two additional analyses of the hazard rate with stratification with respect to age at diagnosis and estrogen receptor status were performed in the Supplementary Materials. The results suggest that the shift in mortality around year 1996 could correspond to the introduction of hormone-blocking therapy.

7 Conclusion

In this article, we have introduced a new estimation method to deal with age-period-cohort analysis. This model assumes no specific structure of the effects of age and cohort and the hazard rate is directly estimated without estimating the effects. In order to take into account possible overfitting issues, a penalty is used on the likelihood to enforce similar consecutive values of the hazard to be equal. Two different types of penalty terms were introduced. One leads to a ridge type regularization while the other leads to a L_0 regularization. Different selection methods of the penalty parameter were also introduced. To our knowledge, a segmented estimation model of this kind has never been introduced in this context.

Using simulated data, it has been shown that the cross validated ridge estimator and the $EBIC_0$ adaptive ridge estimator perform the best in terms of mean squared error. The cross validation criterion was shown to provide the best fit of the hazard rate, but its high computational cost makes it non-competitive. In this context, this modified BIC criterion comes out as a powerful tool to select the best bias-variance tradeoff.

The method was successfully applied to data of survival after breast cancer provided by the SEER program. The segmented estimate of the hazard rate displays important information about the shift in mortality after being diagnosed of breast cancer in the United States in the mid-1990s.

Our method could be directly extended to a different discretization of the age-period-cohort plane, such as $1 \times 1 \times 1$ -year triangles that are represented in dark gray in Figure 1 (see Section 3 of Carstensen [5], for an example of this discretization). Another extension would be to consider other types of penalizations. Instead of estimating a piecewise constant hazard, one could estimate a piecewise linear hazard by penalizing over second order differences of the hazard.

Acknowledgements: The authors are thankful to the National Cancer Institute for providing U.S. mortality data on cancer. We also thank an anonymous reviewer for comments on the adaptive ridge and its link to related methods which significantly improved the quality of this paper.

Author contribution: All the authors have accepted responsibility for the entire content of this submitted manuscript and approved submission.

Research funding: None declared.

Conflict of Interest statement: The authors have declared no conflict of interest.

References

1. Yang Y, Land KC. Age-period-cohort analysis: new models, methods, and empirical applications. Chapman & Hall/CRC Interdisciplinary Statistics; 2013.
2. Osmond C, Gardner MJ. Age, period and cohort models applied to cancer mortality rates. *Stat Med* 1982;1:245–59.
3. Heuer C. Modeling of time trends and interactions in vital rates using restricted regression splines. *Biometrics* 1997;53:161–77.
4. Holford TR. The estimation of age, period and cohort effects for vital rates. *Biometrics* 1983;39:311–24.
5. Carstensen B. Age–period–cohort models for the Lexis diagram. *Stat Med* 2007;26:3018–45.
6. Kuang D, Nielsen B, Nielsen JP. Identification of the age-period-cohort model and the extended chain-ladder model. *Biometrika* 2008;95:979–86.
7. Nielsen B. Apc: an R package for age-period-cohort analysis. *The R Journal* 2015;7:52.
8. Carstensen B, Plummer M, Laara E, Hills M. Epi: a package for statistical analysis in epidemiology; 2017.
9. Plummer M, Carstensen B. Lexis: an R class for epidemiological studies with long-term follow-up. *J Stat Software* 2011;38:1–12.
10. Beran R. Nonparametric regression with randomly censored survival data. Technical Report. Berkeley: University of California; 1981.
11. McKeague IW, Utikal KJ. Identifying nonlinear covariate effects in semimartingale regression models. *Probab Theor Relat Field* 1990;87:1–25.
12. Keiding N. Statistical inference in the Lexis diagram. *Phil Trans Roy Soc Lond A* 1990;332:487–509.
13. Currie ID, Kirkby JG. Smoothing age-period-cohort models with P-splines: a mixed model approach; 2009.
14. Candès EJ, Wakin MB, Boyd SP. Enhancing sparsity by reweighted L1 minimization. *J Fourier Anal Appl* 2008;14:877–905.
15. Chartrand R, Yin W. Iteratively reweighted algorithms for compressive sensing. In: 2008 IEEE International Conference on Acoustics, Speech and Signal Processing; 2008:3869–72 pp.
16. Rippe RCA, Meulman JJ, Eilers PHC. Visualization of genomic changes by segmented smoothing using an L0 penalty. *PLoS One* 2012;7:e38230.
17. Frommlet F, Nuel G. An adaptive ridge procedure for L0 regularization. *PLoS One* 2016;11:e0148620.
18. Bouaziz O, Nuel G. L0 regularization for the estimation of piecewise constant hazard rates in survival analysis. *Appl Math* 2017;08:377–94.
19. Aalen O, Borgan O, Gjessing H. Survival and event history analysis: a process point of view. Springer Science & Business Media; 2008.
20. Ogata Y, Katsura K. Likelihood analysis of spatial inhomogeneity for marked point patterns. *Ann Inst Stat Math* 1988;40:29–39.
21. Foucart S, Lai M-J. Sparsest solutions of underdetermined linear systems via Lq-minimization for $0 < q \leq 1$. *Appl Comput Harmon Anal* 2009;26:395–407.
22. Wipf D, Nagarajan S. Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions. *IEEE J Sel Top Signal Process* 2010;4:317–29.
23. Csardi G, Nepusz T. The igraph software package for complex network research; 2006.
24. Schwarz G. Estimating the dimension of a model. *Ann Stat* 1978;6:461–4.
25. Żak-Szatkowska M, Bogdan M. Modified versions of the Bayesian information criterion for sparse generalized linear models. *Comput Stat Data Anal* 2011;55:2908–24.
26. Chen J, Chen Z. Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* 2008;95:759–71.
27. Akaike H. Information theory and an extension of the maximum likelihood principle. In Selected papers of Hirotugu Akaike. Springer; 1998:199–213 pp.

Supplementary Material: The online version of this article offers supplementary material (<https://doi.org/10.1515/ijb-2019-0003>).