International Journal of Approximate Reasoning 80 (2017) 239–256

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Bayes linear kinematics in a dynamic survival model

Kevin J. Wilson *, Malcolm Farrow

School of Mathematics and Statistics, Newcastle University, Newcastle upon Tyne, UK

ARTICLE INFO

Article history:
Received 27 November 2015
Received in revised form 8 September 2016
Accepted 23 September 2016
Available online 28 September 2016

Keywords:
Bayes linear kinematics
Bayes linear Bayes graphical model
Dynamic model
Piecewise constant hazard
Survival analysis
Time-dependent covariate effects

ABSTRACT

Bayes linear kinematics and Bayes linear Bayes graphical models provide an extension of Bayes linear methods so that full conditional updates may be combined with Bayes linear belief adjustment. In this paper we investigate the application of this approach to survival analysis with time-dependent covariate effects, a more complicated problem than previous applications. We use a piecewise-constant hazard function with a prior in which covariate effects are correlated over time. The need for computationally intensive methods is avoided and the relatively simple structure facilitates interpretation. Our approach eliminates the problem of non-commutativity which was observed in earlier work by Gamerman. We apply the technique to data on survival times for leukemia patients.

1. Introduction

A Bayes linear analysis [6] differs from a full Bayesian analysis in that only first and second order moments are specified in the prior. Posterior (termed *adjusted*) moments are then calculated when data are observed. The introduction of Bayes linear kinematics and Bayes linear Bayes models [5] extends Bayes linear methods to allow the incorporation of observations of types which are not readily accommodated in a straightforward Bayes linear analysis. For example, beliefs about certain unknown quantities might be updated by full conditional Bayesian inference when observations are made on conditionally Poisson or binomial variables and then information can be propagated between these unknowns, or to other unknowns, via a Bayes linear belief structure. This approach avoids the need for computationally intensive methods such as Markov chain Monte Carlo which are often required in standard Bayesian analyses. Computational time in evaluating posterior distributions can be an important issue in areas such as design of experiments [19], clinical decision rules or evaluation of diagnostic tests. Such analyses may require the repeated evaluation of posterior distributions given large numbers of simulated data sets. In such cases, the Bayes linear kinematic method may provide an effective emulator [12].

Wilson and Farrow [26] introduced the use of a link function to map the range of an unknown, such as the mean of a Poisson distribution, onto the whole real line and improve the linearity of the relationships represented by the Bayes linear structure. In this paper we show how Bayes linear kinematics may be applied to a more complicated problem, specifically in the analysis of survival data, and that this brings appealing advantages over standard techniques. Our analysis uses death and censoring times, in contrast to the relatively simple actuarial methods developed in [26]. We use a piecewise constant hazards model with temporally-dependent hazard priors. We combine fully Bayesian conjugate updating for individuals in intervals and Bayes linear kinematic updating to propagate changes in belief to other individuals and intervals. Our

* Corresponding author.
E-mail addresses: kevin.wilson@newcastle.ac.uk (K.J. Wilson), malcolm.farrow@newcastle.ac.uk (M. Farrow).

http://dx.doi.org/10.1016/j.ijar.2016.09.010
0888-613X/

model is related to that of Gamerman [3] but, using the Bayes linear kinematic approach, we overcome the problem of non-commutativity of updates observed by Gamerman.

We consider Bayesian analysis from a subjectivist perspective [4,16]. Therefore we give attention to the appropriate specification of prior beliefs.

The remainder of the paper is structured as follows. Section 2 gives an overview of proportional hazards models and the piecewise constant hazards model and reviews the model of Gamerman [3]. In Section 3 we give a brief introduction to the results of Goldstein and Shaw [5]. In Section 4 we describe our Bayes linear kinematic solution to the survival problem in four stages; the guide relationship, system evolution, use of Bayes linear kinematics and calculation of the expectations and variances. The usefulness of the approach is illustrated with an example involving survival times of leukemia patients in the North-West of England in Section 5. Some conclusions and areas for further work are presented in Section 7.

2. Survival analysis

2.1. Introduction

In this paper we investigate the application of Bayes linear kinematics and Bayes linear Bayes models in survival analysis, specifically a proportional hazards model with piecewise constant hazards. Survival analysis is concerned with modelling the time elapsed, known as the survival time, until some event occurs. For convenience we shall refer to the event as "death".

The survival time t of an individual is a realisation of a random variable T. Associated with T is a survival function $S(t) = \Pr(T \geq t)$, a probability density function $f(t)$ and a hazard function $h(t) = f(t)/S(t)$. Censoring of observations is a common feature of survival data. In right censoring all that is known is that $t > c$ for some value c, in left censoring this condition is $t < c$ and in interval censoring $c_1 < t < c_2$, for some values c_1, c_2. In this paper we consider only right censoring, which is the most common type, and assume that the censoring is non-informative. That is, the survival time T is independent of the mechanism which causes an observation to be censored. Further information on Bayesian survival analysis can be found in [15] and [10].

2.2. Proportional hazards models

Suppose we have individuals $i = 1, \ldots, p$ and individual i has covariate values $x_i = (x_{i,0}, x_{i,1}, \ldots, x_{i,q})'$ where, typically, $x_{i,0} \equiv 1$. Associated with individual i is a hazard function $h_i(t)$. A popular and appealing way to relate the covariate values to the survival distribution for an individual is to make the proportional hazards assumption [1]. Then we can write $h_i(t) = \phi_i h_0(t)$, where ϕ_i is a constant with respect to time and $h_0(t)$ is a baseline hazard function. We can relate an individual's hazard function to x_i, the individual's covariate vector, by setting

$$\phi_i = \exp\left(\sum_{k=1}^{q} x_{i,k}\beta_k\right), \tag{1}$$

for some parameters $\beta_1, \ldots, \beta_q$, which, in a simple proportional hazards model, remain constant over time.

2.3. Piecewise constant hazards model

We might be unwilling to assume a particular form for the baseline hazard function $h_0(t)$. A simple and much investigated way to relax this assumption is to use a piecewise constant hazards model (e.g. [10]). Time is partitioned into disjoint intervals. In each interval a constant hazard is specified but the hazards are allowed to vary from interval to interval.

Furthermore we may wish to allow the effects of the covariates, represented by the coefficients $\beta_1, \ldots, \beta_q$, to vary from one time interval to another. This has led to the development of dynamic survival models in which the coefficients can vary over time [17]. We shall consider a dynamic model $h_i(t) = \exp\{x_i'\beta(t)\}$, where $x_i' = (1, x_{i,1}, \ldots, x_{i,q})$ and $\beta'(t) = (\beta_0(t), \beta_1(t), \ldots, \beta_q(t))$ with $\beta_0(t) = \log\{h_0(t)\}$, so that we can model changes in the effects of the covariates over time. The static model in (1) is then a special case of this more general model.

We choose fixed time points $\tau_0, \tau_1, \ldots, \tau_r$ such that $\tau_0 = 0$ and $\tau_r \to \infty$. This partitions time into intervals. We say that the jth interval is $R_j = [\tau_{j-1}, \tau_j)$. Then, for $\tau_{j-1} \leq t < \tau_j$, the baseline hazard is $h_0(t) = \lambda_{0,j}$ and the hazard function for individual i is $h_i(t) = \lambda_{i,j} = \phi_{i,j}\lambda_{0,j} = \exp(\eta_{i,j})$, where $\eta_{i,j} = x_i'\beta_j$ is the *linear predictor* and $\beta_j = (\beta_{j,0}, \ldots, \beta_{j,q})'$. That is, the hazard for each individual remains constant through each of the time intervals. The integrated hazard $H_i(t) = \int_0^t h_i(u)du$ is then

$$H_i(t) = \sum_{k:\tau_k < t} \lambda_{i,k}(\tau_k - \tau_{k-1}) + \lambda_{i,j}(t - \tau_{j-1}),$$

for $k = 1, \ldots, j-1$.

If we condition on $T \geq \tau_j$ then we obtain the conditional survival function and conditional probability density function for individual i at time t. These are, for $\tau_{j-1} \leq t < \tau_j$,