

A Multistate Model for Bivariate Interval-Censored Failure Time Data

RICHARD J COOK

*Department of Statistics and Actuarial Science,
University of Waterloo, Waterloo, ON, N2L 3G1, Canada
E-mail: rjcook@uwaterloo.ca,*

LEILEI ZENG

*Department of Statistics and Actuarial Science,,
Simon Fraser University, Burnaby, BC, V5A 1S6, Canada*

KER-AI LEE

*Department of Statistics and Actuarial Science,
University of Waterloo, Waterloo, ON, N2L 3G1, Canada*

Summary

Interval-censored life-history data arise when the events of interest are only detectable at periodic assessments. When interest lies in the occurrence of two such events, bivariate-interval censored event time data are obtained. We describe how to fit a four-state Markov model useful for characterizing the association between two interval-censored event times when the assessment times for the two events may be generated by different inspection processes. The approach treats the two events symmetrically and enables one to fit multiplicative intensity models that give estimates of covariate effects as well as relative risks characterizing the association between the two events. An expectation-maximization (EM) algorithm is described for estimation in which the maximization step can be carried out with standard software. The method is illustrated by application to data from a trial of HIV patients where the events are the onset of viral shedding in the blood and urine among individuals infected with cytomegalovirus.

Keywords: Bivariate failure time; EM algorithm; Interval censoring; Markov model

This is the peer reviewed version of the following article: Cook, R. J., Zeng, L. and Lee, K.-A. (2008), A Multistate Model for Bivariate Interval-Censored Failure Time Data. *Biometrics*, 64: 1100-1109. doi: 10.1111/j.1541-0420.2007.00978.x, which has been published in final form at <http://onlinelibrary.wiley.com/doi/10.1111/j.1541-0420.2007.00978.x/full>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving: <http://olabout.wiley.com/WileyCDA/Section/id-820227.html#terms>.

1 INTRODUCTION

Studies of life history data are often designed to collect information on the times of different kinds of events. When individuals are examined continuously over a period of follow-up the precise event times are either known or are right censored. When the events are asymptomatic, more careful assessment of individuals is required to detect their occurrence, often through radiological examination, serological tests, or thorough clinical examination. Examples include the development of fractures in patients with osteoporosis detected by radiographic examinations (Riggs et al. (1990)), the occurrence of skeletal metastases in cancer patients, which are detected by bone scans (Chen et al. (2005)), the development of serological responses in orthopedic surgery (Warkentin et al. (2005)), and the occurrence of viral shedding in patients with cytomegalovirus (CMV) infection (Betensky and Finkelstein (1999b); Goggins and Finkelstein (2000)). Interval-censored data typically result from such assessment schemes; see Sun (2006) for an excellent survey of recent developments in this area.

Sometimes interest lies in the occurrence of more than one event subject to interval censoring. For the AIDS Clinical Trials Group study (ACTG 181), laboratory tests were performed at periodic times to test for the presence of CMV in the blood and at possibly different times to test for the presence of CMV in the urine (called CMV shedding in the blood and urine, respectively). Urine samples of patients were drawn every 4 weeks whereas blood samples were drawn every 12 weeks at the clinic visit. In these studies, the time to the event lies between the last visit of negative test and the first visit of positive test, which results in discrete interval-censored failure time data for each event processes. In addition, because one process is inspected more frequently than the other process for the patients, the intervals into which each of the failures are censored are of varying lengths and may overlap.

With bivariate interval-censored data, the most appropriate way to proceed with the analysis depends on the questions of primary scientific importance. In some cases the association between the two events is a nuisance and it must be dealt with simply to ensure a valid basis for simultaneous inference regarding different features of the event process. Goggins and Finkelstein (2000) and Kim and Xue (2002), for example, describe marginal methods for the analysis of bivariate interval-censored failure time data in which robust variance estimates are derived to accommodate a correlation in the event times; this idea is similar in spirit to the use of robust sandwich-type variance estimates in a working independence model for clustered data (Liang and Zeger (1986)) or multivariate failure time data (Wei et al. (1989)). In this framework, interest lies primarily in the regression coefficients for the marginal discrete time proportional hazards regression models. Fully specified models can also be constructed by combining marginal distributions of any specified form with a suitable copula to form a valid and flexible multivariate distribution (Joe (1997); Hougaard (2000)). An attractive feature of copula models is that the margins do not depend on the choice of the dependency structure so that they can be modeled and estimated separately. It is due to this feature, in part, that two-stage estimation procedures have been proposed when marginal analyses are considerably more straightforward (Shih and Louis (1995)).

In other settings interest may lie primarily in the association between two or more events. Betensky and Finkelstein (1999b) develop a nonparametric estimate of a bivariate survival distribution in the setting of a discrete time model but do not consider the impact of covariates. An extension of Kendall's coefficient of concordance was proposed by Betensky and Finkelstein (1999a) for a careful assessment of the association between failure times. Other frameworks for multivariate or clustered interval-censored failure time data are based on frailty models (Bellamy et al. (2004)). Models which furnish relative risks to characterize the dependence can also be formed but to date this approach has been used only for the case of an interval-censored time-dependent covariate and a right-censored event time (Goggins et al. (1999)).

We propose a method for the analyses of bivariate interval censored failure time data based on a multistate Markov formulation. The framework is the same as that discussed in Andersen et al. (1993)

for modeling the association between right-censored bivariate failure time data. In the framework of a multistate model, the problem has similarities to the problem of fitting Markov models to panel data considered by Kalbfleisch and Lawless (1985) and Kalbfleisch and Lawless (1989) who developed a Fisher-scoring algorithm for estimation. A unique feature of the current setting, however, is that the assessment times for the two different events are not necessarily the same, so the censoring times are different for different kinds of transitions. We propose an expectation maximization (EM) algorithm for estimation in a similar spirit to Lindsey and Ryan (1993) and Lindsey and Ryan (1998). The ability to deal with different inspection time processes is an important feature of our approach. Moreover, our method enables one to understand the associations between the two event times and helps to answer the questions regarding which type of event tends to occur first, and how the risk of one event alters risk of the other. Moreover, it provides estimates of meaningful marginal quantities from the multistate model as we illustrate in the example.

The remainder of this article is organized as follows. The multistate formulation is introduced in Section 2. In Section 3 we discuss estimation of quantities of common interest, including the marginal survivor function. The proposed methods are applied to the ACTG 181 in Section 4. General remarks are made in Section 5.

2 A MULTISTATE APPROACH FOR BIVARIATE INTERVAL-CENSORED DATA

2.1 NOTATION AND MODEL FORMULATION

For expository purposes, we discuss the data, model formulation, and algorithm for a single individual. Let T_j denote the time to occurrence of the type j event with corresponding marginal survivor function $\mathcal{F}_j(t) = P(T_j \geq t)$, $j = 1, 2$. We consider the setting where the assessment times for the two different types of events are not necessarily the same and so let $0 = v_{j0} < v_{j1} < \dots < v_{jr_j}$ denote the r_j assessment times for type j events, $j = 1, 2$, which we assume satisfies the conditions of Grüger et al. (1991). At each assessment time for type j events we observe realizations of $Y_{jk} = I(T_j \leq v_{jk})$, $k = 1, \dots, r_j$, $j = 1, 2$. If $Y_{jk} = 1$ for some k , then the censoring interval for T_j is then $C_j = (L_j, R_j]$ where

$$L_j = \max_{k: Y_{jk}=0} (v_{jk}) \quad \text{and} \quad R_j = \min_{k: Y_{jk}=1} (v_{jk}) ;$$

if $Y_{jk} = 0$ for all $k = 1, \dots, r_j$, then $L_j = v_{jr_j}$ and $R_j = \infty$, $j = 1, 2$. We know, therefore, that $T \in \mathcal{C}$ where $T = (T_1, T_2)'$ and $\mathcal{C} = C_1 \times C_2$.

Consider the four-state model depicted in Figure 1 in which the states indicate the status with respect to the occurrence of the two types of events at a given time. If an individual is in state 0 at time t , then neither event has occurred (e.g., $(T_1 \geq t, T_2 \geq t)$), if they are in state 1, only the type 1 event has occurred (e.g., $(T_1 < t, T_2 \geq t)$), if they are in state 2, only the type 2 event has occurred (e.g., $(T_1 \geq t, T_2 < t)$), and if they are in state 3, both types of events have occurred (e.g., $(T_1 < t, T_2 < t)$). Let $\{Z(s), 0 < s\}$ denote the stochastic process characterizing the transitions through the states in Figure 1, where $Z(0) = 0$. We may then write $\mathcal{F}_1(t) = P(Z(t) \in \{0, 2\} | Z(0) = 0)$ and $\mathcal{F}_2(t) = P(Z(t) \in \{0, 1\} | Z(0) = 0)$.

Let $N_{\ell k}(t)$ count the number of $\ell \rightarrow k$ transitions over $(0, t]$, $dN_{\ell k}(t) = 1$ if a $\ell \rightarrow k$ transition occurs at time t , and $\mathbf{N}(t) = (N_{01}(t), N_{02}(t), N_{13}(t), N_{23}(t))'$. The intensity for $\ell \rightarrow k$ transitions are defined as

$$\lambda_{\ell k}(t | H(t)) = \lim_{\Delta t \rightarrow 0} \frac{P(N_{\ell k}((t + \Delta t)^-) - N_{\ell k}(t^-) = 1 | Z(t^-) = \ell, H(t))}{\Delta t}, \quad (2.1)$$

where $H(t) = \{Z(s), 0 < s < t\}$ denotes the history of the process at t . For Markov models $\lambda_{\ell k}(t | H(t)) = \lambda_{\ell k}(t)$ because it is sufficient to condition on the current state occupied to characterize the instantaneous probability of a transition.

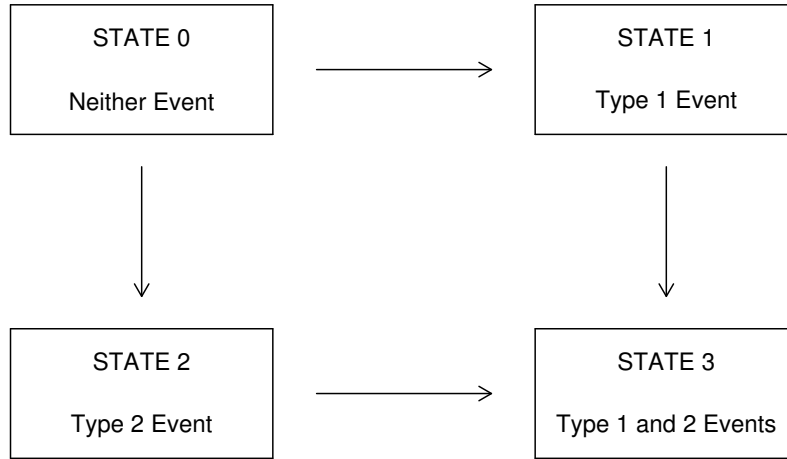


Figure 1: A four-state model for bivariate failure time data.

If it is of interest to determine how the risk of one event changes with the occurrence of the other event, one can compare plots of cumulative intensity functions. For example, if $\Lambda_{23}(t) = \int_0^t \lambda_{23}(u)du$ versus t is much steeper than $\Lambda_{01}(t) = \int_0^t \lambda_{01}(u)du$ versus t , then one may conclude that the intensity for type 1 events is increased after the occurrence of a type 2 event. Similarly, comparing the plots of $\Lambda_{02}(t)$ with $\Lambda_{13}(t)$ gives us idea of whether the risk of type 2 event alters after the occurrence of type 1 event. A more parsimonious summary of this kind of dependence is possible by specifying multiplicative models for transition intensities of the form

$$\begin{aligned} \lambda_{23}(t) &= \lambda_{01}(t) \exp(\alpha_1) \\ \lambda_{13}(t) &= \lambda_{02}(t) \exp(\alpha_2), \end{aligned} \tag{2.2}$$

where $\exp(\alpha_1)$ and $\exp(\alpha_2)$ are relative risks characterizing how the risk of one type event changes with the occurrence of the other. Goodness-of-fit assessment can be carried out to check the plausibility of the model in (2.2) compared to the more general model. Moreover, if (2.2) appears to provide an adequate description of the data, then two degree of freedom (d.f.) tests of $H_0 : \alpha_1 = \alpha_2 = 0$ are tests of independence between the two events.

Suppose $\mathbf{x} = (x_1, \dots, x_p)'$ is a $p \times 1$ covariate vector associated with the transition intensities. The effects of covariates can be modeled in a number of ways, but multiplicative models of the form

$$\lambda_{\ell k}(t|x) = \lambda_{\ell k 0}(t) \exp(\mathbf{x}'\beta_{\ell k}), \tag{2.3}$$

are perhaps the most common, where $\beta_{\ell k}$ is a $p \times 1$ vector of regression coefficients and $\lambda_{\ell k 0}$ is the baseline rate function for the $\ell \rightarrow k$ transition. The covariate vectors may of course be different for the different transition intensities but we assume they are common here for notational convenience only. If we consider introducing the constraint in (2.2) into the model (2.3), then we can consider the resulting model as comprising coupled proportional intensity models for the events of type 1 and 2. For the intensity for type 1 events, for example, $N_{02}(t)$ is an internal time-dependent covariate indicating whether the type 2 event has occurred; the time at which it becomes one is of course interval censored.

For the baseline intensity functions we adopt flexible piecewise constant forms that offer robustness to misspecification compared to more strongly parametric models, and convenience of estimation as demonstrated in Section 2.2. Let $0 = b_{\ell k 0} < b_{\ell k 1} < \dots < b_{\ell k, p_{\ell k}-1} < b_{\ell k, p_{\ell k}} = \infty$ denote the specified breakpoints for the intensity for $\ell \rightarrow k$ transitions, and $\mathcal{B}_{\ell k h} = (b_{\ell k, h-1}, b_{\ell k h}]$ be the h th interval, $h = 1, \dots, p_{\ell k}$. For piecewise constant transition intensity functions we set $\lambda_{\ell k 0}(t) = \lambda_{\ell k h}$ for $t \in \mathcal{B}_{\ell k h}$, $h = 1, \dots, p_{\ell k}$ and let $\lambda_{\ell k} = (\lambda_{\ell k 1}, \dots, \lambda_{\ell k p_{\ell k}})'$ be the $p_{\ell k} \times 1$ vector of parameters associated with baseline $\ell \rightarrow k$ transition intensity function. If we let $d_{\ell k h}(t) = I(b_{\ell k, h-1} < t \leq b_{\ell k h})$, the transition intensity function can be written as

$$\lambda_{\ell k}(t|x) = \prod_{h=1}^{p_{\ell k}} \lambda_{\ell k h}^{d_{\ell k h}(t)} \exp(x' \beta_{\ell k}), \quad (2.4)$$

and the integrated transition intensity function is

$$\Lambda_{\ell k}(s, t|x) = \int_s^t \lambda_{\ell k}(u|x) du = \sum_{h=1}^{p_{\ell k}} u_{\ell k h}(s, t) \lambda_{\ell k h} \exp(x' \beta_{\ell k}), \quad (2.5)$$

where $u_{\ell k h}(s, t) = \max(\min(b_{\ell k h}, t) - \max(b_{\ell k, h-1}, s), 0)$ equals the length of the intersection of interval $(s, t]$ with interval $\mathcal{B}_{\ell k h}$. Under most circumstances, piecewise constant intensity models with relatively few pieces are sufficient for analysis (Lawless and Zhan (1998)). Let $\lambda = (\lambda'_{01}, \lambda'_{02}, \lambda'_{13}, \lambda'_{23})'$, $\beta = (\beta'_{01}, \beta'_{02}, \beta'_{13}, \beta'_{23})'$, and $\theta = (\lambda', \beta)'$ contain all parameters of the model.

2.2 AN EM ALGORITHM

We describe the following EM algorithm (Dempster et al. (1977)) in the absence of covariates, but remark that introduction of covariates via (2.2) is straightforward. Given the multistate formulation in Section 2.1, we consider the ‘‘complete’’ data for this EM algorithm as including the transition times in Figure 1. The complete data likelihood is then

$$L_C(\theta) \propto [\lambda_{01}(t_1) \exp\{-\Lambda_{01}(0, t_1) - \Lambda_{02}(0, t_1)\} \cdot \lambda_{13}(t_2) \exp\{-\Lambda_{13}(t_1, t_2)\}]^{I(t_1 < t_2)} \\ \times [\lambda_{02}(t_2) \exp\{-\Lambda_{01}(0, t_2) - \Lambda_{02}(0, t_2)\} \cdot \lambda_{23}(t_1) \exp\{-\Lambda_{23}(t_2, t_1)\}]^{1-I(t_1 < t_2)},$$

where the indicator function in the exponent determines the form of the contribution. With the piecewise constant model for the transition intensities, the complete data log likelihood may be written as

$$\ell_C(\theta) = I(t_1 < t_2) \cdot \ell_C^1(\theta) + [1 - I(t_1 < t_2)] \cdot \ell_C^2(\theta), \quad (2.6)$$

where

$$\ell_C^1(\theta) = \sum_{h=1}^{p_{01}} d_{01h}(t_1) \log \lambda_{01h} - \sum_{k=1}^2 \sum_{h=1}^{p_{0k}} u_{0kh}(0, t_1) \lambda_{0kh} + \sum_{h=1}^{p_{13}} d_{13h}(t_2) \log \lambda_{13h} - \sum_{h=1}^{p_{13}} u_{13h}(t_1, t_2) \lambda_{13h}$$

and

$$\ell_C^2(\theta) = \sum_{h=1}^{p_{02}} d_{02h}(t_2) \log \lambda_{02h} - \sum_{k=1}^2 \sum_{h=1}^{p_{0k}} u_{0kh}(0, t_2) \lambda_{0kh} + \sum_{h=1}^{p_{23}} d_{23h}(t_1) \log \lambda_{23h} - \sum_{h=1}^{p_{23}} u_{23h}(t_2, t_1) \lambda_{23h},$$

where $\ell_C^k(\theta)$ has the form of the log likelihood contribution for the path $0 \rightarrow k \rightarrow 3$, $k = 1, 2$. Under the bivariate interval-censored setting, the failure times (t_1, t_2) are not observed but they are known to be contained in the rectangle $\mathcal{C} = [L_1, R_1] \times [L_2, R_2]$ as shown in Figure 2. Treating t_1 and t_2 as

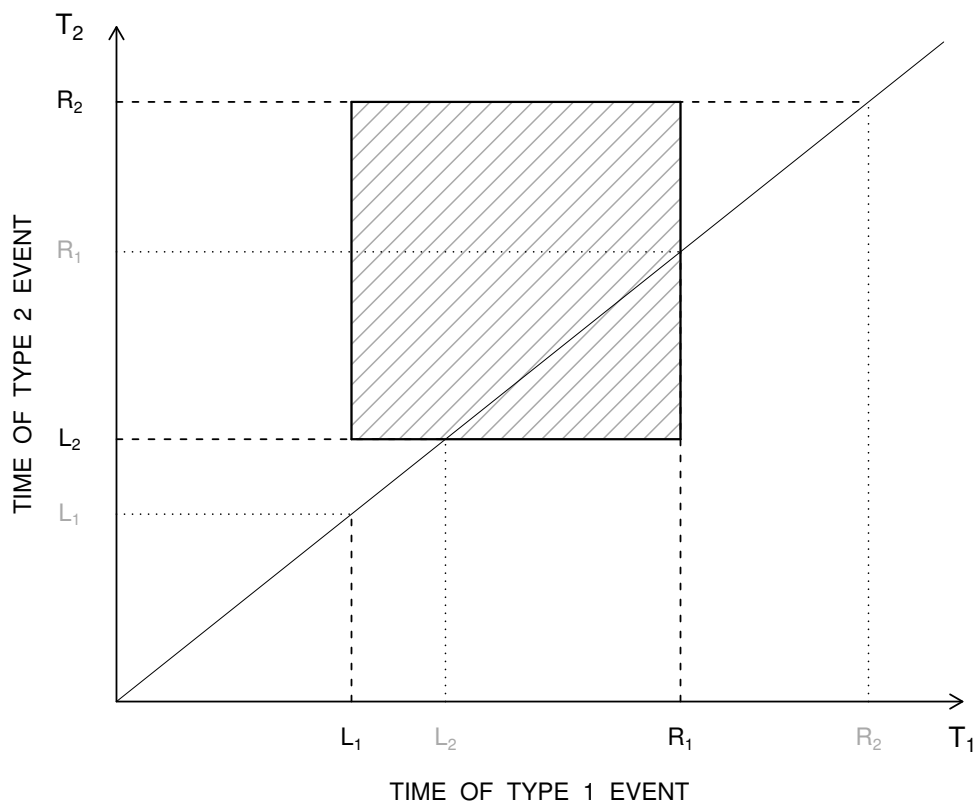


Figure 2: The censoring region for interval censored observation of $(T_1, T_2)'$.

“missing” data, the EM algorithm is used for numerical maximization of the observed data likelihood and estimation of θ .

The E-step involves taking the expectation of the complete data log likelihood (2.6) with respect to $\mathbf{t} = (t_1, t_2)$ given the data (i.e., that $\mathbf{t} \in \mathcal{C}$). If the censoring region lies entirely above or below the line $T_1 = T_2$ then the path ($0 \rightarrow 1 \rightarrow 3$ or $0 \rightarrow 2 \rightarrow 3$) is known, but otherwise it is not. We define $\mathcal{A}^1 = \{(t_1, t_2) \mid \mathbf{t} \in \mathcal{C}, t_1 < t_2\}$ as the full set of possible times within the censoring region \mathcal{C} corresponding to the path $0 \rightarrow 1 \rightarrow 3$, the superscript 1 on \mathcal{A}^1 is a reminder that the first event to occur is a type 1 event. Similarly, we define $\mathcal{A}^2 = \{(t_1, t_2) \mid \mathbf{t} \in \mathcal{C}, t_1 > t_2\}$ for the path $0 \rightarrow 2 \rightarrow 3$, and $\mathcal{C} = \mathcal{A}^1 \cup \mathcal{A}^2$. Graphically, \mathcal{A}^1 and \mathcal{A}^2 represent the regions in \mathcal{C} of Figure 2 above and below the 45 degree line, respectively. At the r th iteration of the EM, the expectation of the complete data log likelihood (2.6), denoted $Q(\theta; \hat{\theta}^{r-1}) = E\{\ell_C(\theta) \mid \mathbf{t} \in \mathcal{C}; \hat{\theta}^{r-1}\}$ can be taken as following fashion

$$Q(\theta; \hat{\theta}^{r-1}) = \sum_{j=1}^2 E\{\ell_C(\theta) \mid \mathbf{t} \in \mathcal{A}^j; \hat{\theta}^{r-1}\} P(\mathbf{t} \in \mathcal{A}^j \mid \mathbf{t} \in \mathcal{C}; \hat{\theta}^{r-1}). \quad (2.7)$$

Here note that when $C_1 \cap C_2 = \emptyset$, the censoring region does not cross the line $T_1 = T_2$, and the path is known. Conditioning on $\mathbf{t} \in \mathcal{A}^j$ in the expectation for (2.7) is therefore redundant because it is clear which term should contribute in the log likelihood (2.6).

From (2.7) one can see that conditional expectations required are relatively easy to evaluate because they may be computed first by conditioning on a particular path in the four-state model. For the following explanation we consider the case with $\mathbf{t} \in \mathcal{A}^1$ and emphasize the associated conditional expectations,

$$\begin{aligned} \delta_{01h}(t_1) &= E\{d_{01h}(t_1) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}, & \delta_{13h}(t_2) &= E\{d_{13h}(t_2) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}, \\ \mu_{0kh}(0, t_1) &= E\{u_{0kh}(0, t_1) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}, & \text{and } \mu_{13h}(t_1, t_2) &= E\{u_{13h}(t_1, t_2) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}, \end{aligned}$$

$k = 1, 2$. To calculate terms $\delta_{01h}(t_1)$ and $\delta_{13h}(t_2)$, note that the region \mathcal{A}^1 can be divided into a set of mutually exclusive sub-regions, within each of which the transition intensities are constant. We define $\mathcal{A}_{hh'}^1 = \{(t_1, t_2) \mid t_1 \in \mathcal{B}_{01h}, t_2 \in \mathcal{B}_{13h'}, \mathbf{t} \in \mathcal{A}^1\}$ as the set of possible *pairs of event times* occurring in intervals \mathcal{B}_{01h} and $\mathcal{B}_{13h'}$, respectively, then $\mathcal{A}^1 = \cup_{hh'} \mathcal{A}_{hh'}^1$. It can be shown that

$$\delta_{01h}(t_1) = P(d_{01h}(t) = 1 \mid \mathbf{t} \in \mathcal{A}^1; \theta) = \sum_{h'=1}^{p_{13}} P(d_{01h}(t_1) = 1, d_{13h'}(t_2) = 1 \mid \mathbf{t} \in \mathcal{A}^1; \theta),$$

where $P(d_{01h}(t_1) = 1, d_{13h'}(t_2) = 1 \mid \mathbf{t} \in \mathcal{A}^1; \theta)$ is obtained as

$$\frac{\int_{\mathcal{A}_{hh'}^1} \lambda_{01h} \exp\{-\Lambda_{01}(0, t_1) - \Lambda_{02}(0, t_1)\} \lambda_{13h'} \exp\{-\Lambda_{13}(t_1, t_2)\} dt_1 dt_2}{\sum_{h=1}^{p_{01}} \sum_{h'=1}^{p_{13}} \int_{\mathcal{A}_{hh'}^1} \lambda_{01h} \exp\{-\Lambda_{01}(0, t_1) - \Lambda_{02}(0, t_1)\} \lambda_{13h'} \exp\{-\Lambda_{13}(t_1, t_2)\} dt_1 dt_2}. \quad (2.8)$$

In a similar manner, $\delta_{13h}(t_2) = P(d_{13h}(t_2) = 1 \mid \mathbf{t} \in \mathcal{A}^1; \theta)$ can be obtained. We likewise define $\mathcal{A}_{hh'}^2 = \{(t_1, t_2) \mid t_2 \in \mathcal{B}_{02h}, t_1 \in \mathcal{B}_{23h'}, \mathbf{t} \in \mathcal{A}^2\}$ to evaluate expectations conditional on $0 \rightarrow 2 \rightarrow 3$ paths and note $\mathcal{A}^2 = \cup_{hh'} \mathcal{A}_{hh'}^2$.

The conditional expected sojourn times $\mu_{0kh}(0, t_1) = E\{u_{0kh}(0, t_1) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}$, $k = 1, 2$, and $\mu_{13h}(t_1, t_2) = E\{u_{13h}(t_1, t_2) \mid \mathbf{t} \in \mathcal{A}^1; \theta\}$ are evaluated as follows. Let $v_1 \leq v_2 \leq v_3 \leq v_4$ represent the temporally ordered members of the set $\{L_1, L_2, R_1, R_2\}$ and the vector $\mathbf{z} = (z(v_1), z(v_2), z(v_3), z(v_4))'$ indicate the state occupied at v_r , $r = 1, 2, 3, 4$. Let \mathcal{Z} denote the set of all possible realizations of the

state vector z consistent with the data ($z(v_1) = 0$ always), and let \mathcal{Z}^1 be the subset consistent with paths with $T_1 < T_2$ (i.e., where $\mathbf{t} \in \mathcal{A}^1$) (likewise, let \mathcal{Z}^2 be the subset consistent with $\mathbf{t} \in \mathcal{A}^2$). It follows that

$$E\{u_{0kh}(0, t_1) \mid \mathbf{t} \in \mathcal{A}^1; \theta\} = \frac{\sum_{z \in \mathcal{Z}^1} \int_{C_k \cap \mathcal{B}_{0kh}} P(z(u) = 0, z(v_1), z(v_2), z(v_3), z(v_4) \mid z(0) = 0) du}{\sum_{z \in \mathcal{Z}^1} \prod_{r=2}^4 P(z(v_r) \mid z(v_{r-1}))},$$

$k = 1, 2$, and

$$E\{u_{13h}(t_1, t_2) \mid \mathbf{t} \in \mathcal{A}^1; \theta\} = \frac{\sum_{z \in \mathcal{Z}^1} \int_{C_2 \cap \mathcal{B}_{13h}} P(z(u) = 1, z(v_1), z(v_2), z(v_3), z(v_4) \mid z(0) = 0) du}{\sum_{z \in \mathcal{Z}^1} \prod_{r=2}^4 P(z(v_r) \mid z(v_{r-1}))}.$$

The conditional probability $P(\mathbf{t} \in \mathcal{A}^1 \mid \mathbf{t} \in \mathcal{C}; \hat{\theta}^{r-1})$ in (2.7) can be obtained by

$$P(T \in \mathcal{A}^1 \mid T \in \mathcal{C}; \theta) = \frac{P(\mathbf{t} \in \mathcal{A}^1; \theta)}{P(\mathbf{t} \in \mathcal{C}; \theta)} = \frac{P(\mathbf{t} \in \mathcal{A}^1; \theta)}{P(\mathbf{t} \in \mathcal{A}^1; \theta) + P(\mathbf{t} \in \mathcal{A}^2; \theta)}, \quad (2.9)$$

where $P(\mathbf{t} \in \mathcal{A}^1; \theta)$ is expressed in the form as given in the denominator of (2.8). When we condition on $\mathbf{t} \in \mathcal{A}^2$, the corresponding δ -terms, μ -terms can be computed in the similar manner. We obtain (2.7) by evaluating these conditional expectations and probabilities at $\hat{\theta}^{(r-1)}$.

The expectation of the complete data log likelihood for the full sample is the sum of terms (2.7) obtained for each individual. Because the resulting function has the form of a sum of Poisson log likelihoods, standard software for generalized linear models can be used to maximize it at each iteration. Covariates can easily be introduced, and we do so in the example. SAS programs for the analyses that follow are available from the website of this journal.

3 ESTIMATING QUANTITIES OF INTEREST

One of the appealing features of Markov models is the convenient way of computing the transition probability matrix. Here, for simplicity, we consider the case where the break points for the different transition intensities are the same, so $b_{\ell kh} = b_h$. For $s < t$, within a particular interval $\mathcal{B}_h = (b_{h-1}, b_h]$ (i.e., $b_{h-1} \leq s < t \leq b_h$), the transition probability matrix is $\mathcal{P}_h(s, t) = [p_{\ell k}(s, t)]$ with elements $p_{\ell k}(s, t) = P(Z(t) = k \mid Z(s) = \ell)$. Note that because the intensities are time homogeneous within \mathcal{B}_h , then

$$\mathcal{P}_h(s, t) = \exp\{(t - s) \times \mathcal{Q}_h\},$$

where \mathcal{Q}_h is the matrix of transition intensities

$$\mathcal{Q}_h = \begin{bmatrix} -\lambda_{01h} - \lambda_{02h} & \lambda_{01h} & \lambda_{02h} & 0 \\ 0 & -\lambda_{13h} & 0 & \lambda_{13h} \\ 0 & 0 & -\lambda_{23h} & \lambda_{23h} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

(Cox and Miller (1977)). If interest lies in computing $P(Z(t) = k \mid Z(0) = \ell)$ where $(0, t]$ spans several of the intervals \mathcal{B}_h , $h = 1, \dots, p$, then

$$\mathcal{P}(0, t) = \prod_{h=1}^{h(t)} \mathcal{P}_h(b_{h-1}, b_h) \times \mathcal{P}_{h(t)+1}(b_{h(t)}, t) \quad (3.1)$$

where $h(t) = \max_{h < t}(h)$. Having obtained this, one can compute estimates of the marginal survivor functions as

$$\begin{aligned}\widehat{\mathcal{F}}_1(t) &= P(Z(t) = 0 \mid Z(0) = 0; \widehat{\theta}) + P(Z(t) = 2 \mid Z(0) = 0; \widehat{\theta}) \\ \widehat{\mathcal{F}}_2(t) &= P(Z(t) = 0 \mid Z(0) = 0; \widehat{\theta}) + P(Z(t) = 1 \mid Z(0) = 0; \widehat{\theta}).\end{aligned}\tag{3.2}$$

These can be compared with nonparametric estimates for goodness-of-fit assessment. One can also compute estimates of the probability that a type j event will occur before a type $3 - j$ event, $j = 1, 2$. This is given by the cumulative incidence function based on a competing risk analysis for transitions out of state 0. For example,

$$P(\text{type 1 event occurs first}) = \int_0^{\infty} \lambda_{01}(s) \exp(-\Lambda_{01}(0, s) - \Lambda_{02}(0, s)) ds .$$

4 APPLICATION

Here we apply the proposed methods for the analysis of bivariate interval-censored event times to data from an AIDS study ACTG 181 discussed in Betensky and Finkelstein (1999b) and Goggins and Finkelstein (2000); we refer readers to these articles for further details on this study. One purpose of the ACTG 181 study was to determine the natural history of the opportunistic infection CMV in an HIV-infected individual. Infectious CMV may be shed in the body fluids of an infected person, and thus may be found in urine and blood. Testing for the presence of CMV virus in blood and urine is important for HIV-infected patients because shedding of the CMV virus is a precursor for the development of active CMV disease.

We consider data from 204 patients who had one or more follow-up clinic visits with known dates, at which urine and/or blood samples were drawn to test for the presence of CMV virus. Urine samples were to be taken every 4 weeks whereas blood samples were to be taken every 12 weeks. There was considerable variability in the actual visit dates due to compliance and scheduling problems, and missed visits.

The time origin in this analysis was taken to be the time of the first assessment because that represents the point at which detailed follow-up began. This is consistent with what is done in most clinical trials where the time origin is the date of recruitment to the trial even though there may not be a biological rationale for this specification. There were 7 individuals whose blood test at $v_{j0} = 0$ was positive, and for these individuals we added one day to their first inspection time (assuming they were negative at time zero); those 49 individuals testing positive at the first urine sample were treated similarly. Figure 3 shows a plot of the censoring regions for a sample of subjects in the dataset, and the frequency of the different censoring patterns are summarized in Table 1. Note from Table 1 that of the 204 individuals under study, 67 had an interval-censored time to shedding in the urine, and 23 had an interval-censored time to shedding in the blood; a number of shedding times are left censored but more are right censored. It is believed that times to CMV shedding in urine and blood are correlated, however, many patients had positive urine tests and negative blood tests at the same time, and a few had positive blood and negative urine tests.

As mentioned in Section 1, Goggins and Finkelstein (2000) developed methods for marginal analysis which involved fitting discrete time proportional hazards models to estimate covariate effects on each event time under a working independence assumption. Valid inferences were assured by use of a robust variance estimate. Specifically, they considered a model with a common effect of CD4 cell

count (CD4 cells/ $\mu\ell$ < 75 vs. ≥ 75) and found a strong significant effect ($\beta = 0.97$, $SE = 0.197$) so that patients with a baseline CD4 count less than 75 had a 2.64 fold increase in risk of shedding (95% CI: 1.79, 3.88). Estimates of the CD4 variable effect on time to shedding in the urine and blood by separate analyses (Finkelstein (1986)) gave estimated regression coefficients 0.88 ($SE = 0.20$) and 1.09 ($SE = 0.30$), respectively.




Figure 3: Scatter plot of selected censoring intervals for the times to CMV shedding in blood and urine.

The methods we propose can help address questions relating to which type of shedding tends to occur first, as well as how one type of shedding alters the risk of the other. We let $\lambda_{01}(t)$ indicate the intensity of the onset of CMV shedding in the blood at time t given no previous CMV shedding, and $\lambda_{23}(t)$ be the intensity of CMV blood shedding at time t given previous CMV urine shedding. Accordingly, $\lambda_{02}(t)$ is the intensity of the onset of CMV shedding in the urine given no previous CMV shedding, and $\lambda_{13}(t)$ is the intensity of CMV shedding in the urine given prior CMV blood shedding.

Friedman (1982) and Lindsey and Ryan (1998) note that it is advisable to choose breakpoints so that there are roughly the same expected number of events in each interval, which ensures that the pieces are estimated with comparable precision. This can be difficult to achieve with the multistate piecewise constant model, but inspection of the nonparametric estimates of the marginal survivor functions gives some guidance; we consider a multistate model with no covariates and separate piecewise intensity functions for all transitions with breakpoints at 2 and 12 months. These correspond

Table 1: Number of censored observations of each type

		Blood			Total
		Left	Interval	Right	
Urine	Left	5	10	34	49
	Interval	2	13	52	67
	Right	0	0	88	88
	Total	7	23	174	204

roughly to the tertiles of the marginal distributions, conditional on the respective event time occurring by 24 months. Other breakpoints were considered and gave similar results.

Next we consider two types of regression model for the transition intensities. One involves the multiplicative assumption of (2.2), and the second has a shared baseline intensity but with the introduction of an external covariate x_i such that

$$\lambda_{\ell k}(t|x_i) = \lambda_{\ell k 0}(t) \exp(x_i \beta_{\ell k}) .$$

The covariate considered was the baseline CD4 cell count at the time of recruitment to ACTG 181.

The estimates, standard errors and log likelihoods for the analyses based on models for the transition intensities with a general and multiplicative form are reported in Table 2. Standard errors could be obtained based on the approaches of Louis (1982) or Meng and Rubin (1991), or nonparametric or parametric bootstrap, but in this application they were obtained by numerical differentiation of the observed data log likelihood. Recall that if $v_0 = 0$ then $z(v_0) = 0$, and so the observed data log likelihood for a particular individual is

$$\ell(\theta) = \log \left[\sum_{z \in \mathcal{Z}} \prod_{k=1}^4 P(z(v_k)|z(v_{k-1})) \right] , \quad (4.1)$$

where $z = (z(v_1), z(v_2), z(v_3), z(v_4))'$ is defined in the paragraph following (2.8) and the summation in (4.1) is over all possible realizations of z consistent with the data. Although the observed data likelihood is unappealing to use for parameter estimation, upon obtaining parameter estimates via the EM algorithm of Section 2, they can be plugged into (4.1) and finite differencing algorithms can be applied to obtain a suitable Hessian matrix to get the observed information matrix (e.g., Jamshidian and Jennrich (2000)). For this analysis we used the subroutine NLPFDD in PROC IML in SAS version 9.1.3 on a SunOS 5.8 platform.

A likelihood ratio test comparing the general transition intensity model with the multiplicative transition intensity model (no covariates) gives $p = 0.057$, which although not significant at the 5% level gives some cause for concern about the multiplicative model. Figure 4 shows plots the nonparametric maximum likelihood estimate (MLE) of the marginal survivor functions for shedding in the blood and urine (Turnbull (1976)). These were obtained by disregarding any information on the occurrence of the other type of event when estimating the cumulative distribution function of one event. Plots of the estimates of the marginal survival distributions obtained from the bivariate nonparametric estimate of Gentleman and Vandal (2001, 2002) are also given. Finally, estimates based on (3.2) from the general Markov model and the multiplicative Markov model are overlaid. The good agreement between the nonparametric estimate and both piecewise Markov estimates gives reassurance that the multiplicative model is satisfactory for the purpose of estimating these functions. The marginal estimate from the nonparametric bivariate estimate is in close agreement for the time to

shedding in the urine, but is somewhat smaller than the other estimates for the time to shedding in the blood.

Table 2: Estimates from analysis of data on CMV shedding in blood and urine

Parameter	Interval in months	General		Multiplicative models			
		No covariate		No covariate		One covariate	
		Estimate	SE	Estimate	SE	Estimate	SE
λ_{01}	[0, 2)	0.032	0.012	0.029	0.010	0.071	0.034
	[2, 12)	0.003	0.002	0.005	0.002	0.013	0.008
	[12, ∞)	0.011	0.008	0.008	0.004	0.027	0.021
λ_{02}	[0, 2)	0.228	0.029	0.236	0.029	0.303	0.056
	[2, 12)	0.037	0.007	0.033	0.006	0.042	0.011
	[12, ∞)	0.010	0.008	0.023	0.010	0.041	0.022
λ_{13}	[0, 2)	4.983	4.102				
	[2, 12)	0.143	0.093				
	[12, ∞)	2.518	2.340				
λ_{23}	[0, 2)	0.052	0.031				
	[2, 12)	0.012	0.005				
	[12, ∞)	0.015	0.008				
α_1				0.771	0.442	0.204	0.820
α_2				2.540	0.413	5.904	0.985
β_{01}^\dagger						-0.825	0.406
β_{02}^\dagger						-0.285	0.113
β_{23}^\dagger						-0.926	0.566
β_{13}^\dagger						-2.975	1.022
	log-likelihood	-624.964		-629.552		-478.385	

[†] covariate is baseline CD4 cell count; coefficient expressed per 100 unit difference in CD4 count.

The α_1 coefficient in the multiplicative model ignoring CD4 cell count corresponds to a relative risk (RR) of 2.161 (95% CI: 0.909, 5.137; $p=0.081$) indicating a more than twofold increase in risk of shedding in the blood once shedding begins in the urine; likewise the estimate of α_2 gives RR=12.680 (95% CI: 5.643, 28.491; $p < 0.001$) indicating a highly significant and substantial increase in risk of shedding in the urine following the onset of shedding in the blood. Thus there appears to be a strong association between the two event times.

We next add the baseline CD4 cell count to the multiplicative model to examine its effect on the time to shedding. The dataset is substantially reduced for this analysis because there was only covariate information for 174 individuals, so the log likelihoods between the first two models cannot be compared with the third model. A likelihood ratio test of the significance of CD4 cell count can, however, be carried out by fitting the multiplicative model ignoring the CD4 cell count to the reduced dataset comprised of individuals with complete CD4 cell count information. This gives a likelihood ratio statistic of 32.24 which, based on 4 d.f. gives $p < 0.0001$, suggesting that CD4 cell count is




Figure 4: Estimates of cumulative distribution functions from piecewise constant multistate intensities and marginal nonparametric estimates.

an important prognostic factor. The estimated regression coefficients $\hat{\beta}_{01}$ and $\hat{\beta}_{23}$ give relative risks of 0.438 (95% CI: 0.198, 0.970; $p=0.042$) and 0.396 (95% CI: 0.131, 1.202; $p=0.102$), indicating a substantial reduction in the intensity of shedding in the blood for every 100 unit increase in the baseline CD4 count; whereas only the first of these reaches statistical significance, the risk reductions are comparable in magnitude. In contrast, the estimates $\hat{\beta}_{02}$ and $\hat{\beta}_{13}$ are quite different in magnitude and give relative risks of 0.752 (95% CI: 0.603, 0.938; $p=0.011$) and 0.051 (95% CI: 0.007, 0.379; $p=0.004$), respectively. These suggest that among individuals with no evidence of shedding in the blood there is a modest, but statistically significant, reduction in risk of shedding in the urine with every 100 unit increase in the baseline CD4 cell count. Among individuals experiencing shedding in the blood, there is a much more substantial and highly significant reduction in risk of shedding in the urine with higher CD4 cell counts.

We may think of the above model as coupled proportional hazards models for shedding in the blood and urine, with an internal time-dependent covariate indicating whether shedding has been detected by the other test. For example, let $v_{i1}(t) = 1$ if individual i has evidence of shedding in the blood at time t and $v_{i1}(t) = 0$ otherwise; $v_{i2}(t)$ may be similarly defined for shedding in the urine. Then if x_i is the CD4 covariate, the multistate model for shedding in the blood and urine may be equivalently written as the coupled multiplicative intensity models:

$$\zeta_1(t|v_{i2}(t), x_i) = \zeta_{10}(t) \exp(\alpha_{11}x_i + \alpha_{12}v_{i2}(t) + \alpha_{13}x_iv_{i2}(t)) , \quad (4.2)$$

$$\zeta_2(t|v_{i1}(t), x_i) = \zeta_{20}(t) \exp(\alpha_{21}x_i + \alpha_{22}v_{i1}(t) + \alpha_{23}x_iv_{i1}(t)) , \quad (4.3)$$

where $\zeta_{k0}(t)$ is the baseline intensity of type k events. By fitting a model (4.2-4.3) with $\alpha_{13} = \alpha_{23} = 0$ we can test the hypothesis that the regression coefficients in (2.3) are constrained as $\beta_{01} = \beta_{23}$ and $\beta_{02} = \beta_{13}$. In this case, the exponentiated coefficients can be interpreted simply as relative risks for type 1 and 2 events, irrespective of the status of the complementary event. A likelihood ratio test of this reduced model (not shown) against the full model in Table 2 gives $p = 0.017$ indicating the reduced model provides inferior fit to the data. We therefore conclude that the effect of CD4 cell count on event occurrence changes upon the occurrence of one or more of the events; this is driven by the large difference between $\hat{\beta}_{02}$ and $\hat{\beta}_{13}$.

5 DISCUSSION

There is considerable appeal to the use of multistate models to characterize dependencies between events in life-history data. Part of the appeal of this framework is the ability to characterize associations through regression coefficients of internal time-dependent covariates. This familiar framework may be more easily understood than some of the less familiar measures of association such as Kendall's coefficient of concordance and is particularly appealing when there are only two types of events.

This approach has some appeal because it offers weakly parametric estimates of the bivariate distribution, which can be considered along with the bivariate nonparametric estimate of Betensky and Finkelstein (1999b). It has the added appeal of readily dealing with covariates through the formation of multiplicative intensity models. For estimation, the EM algorithm is appealing because software for fitting proportional hazards exponential regression models can be used. The facility to allow covariate effects to change upon the occurrence of other events is a unique feature as well, for settings with multivariate interval-censored data.

Standard analysis of interval-censored data rests on the assumption that the assessment times are independent of the event times. Betensky and Finkelstein (2002b) discuss identifiability issues and tests for various forms of dependence between failure times and inspection processes, which have a direct bearing on the separate analysis of times to shedding in blood or urine, such as those displayed in

Figure 4. In our multistate formulation, we have assumed that the assessment times for the two events are independent of the transition times of the multistate model in the sense of Grüger et al. (1991). This is a weaker assumption than is required for separate marginal analyses displayed in Figure 4, and is therefore a further advantage of the multistate formulation; that is, the inspection times for one process can be dependent on the observed history of the event times for the other processes (Grüger et al. (1991)).

With univariate interval-censored data, Turnbull (1976) derived the regions of support for the nonparametric estimate, so that the nonparametric estimate is undefined outside of these regions; Betensky and Finkelstein (1999b) discuss a similar algorithm for bivariate interval-censored data. Difficulties can arise in fitting multistate models with piecewise constant intensity functions if breakpoints are chosen such that the risk set for particular transitions is of negligible size. The flexibility of choosing different breakpoints for different transition intensities is appealing, but requires more careful planning. The use of a relatively small number of pieces for each transition is a practical compromise; given the total number of transition intensities, even with a small number of breakpoints, the total number of parameters can still be reasonably large.

Adaptations to some other types of processes are relatively straightforward. In fact, this approach is closely connected with the methods of Lindsey and Ryan (1993), developed for fitting an illness death model to interval-censored data arising in survival-sacrifice experiments. In the multivariate context, with three or more types of events the principles of our method can be adaptive but the number of intensity functions and parameters increases rapidly with the most general model. In such settings, constraints such as those proposed in (2.3) would be particularly convenient to adopt. Goodness of fit can be assessed for a particular model, by fitting more general models that include the particular model as a special case. One can then carry out tests of the adequacy of the null model by testing parameters indexing the expansion; failure to reject the null hypothesis gives some assurance that the particular model is adequate.

Cook and Lawless (2007) consider the use of multistate models for the analysis of right-censored single type and multitype recurrent event data, as well as recurrent event data with terminal events. Chen et al. (2005) consider model fitting for interval-censored multitype recurrent event data where the association is introduced through multivariate random effect distributions or suitably parameterized covariance matrices. In principle, one can adopt multistate models of the sort described here for multitype recurrent events but it may be necessary to parameterize dependencies on the history quite heavily rather than stratifying because of a paucity of information.

The piecewise constant formulation considered here was adopted as a computationally appealing and robust framework for analysis. In this framework, model expansion is relatively straightforward and can be used as a means of carrying out diagnostic checks. The Markov assumption may not be particularly strong for estimating survival functions with right-censored data, because several authors have recently shown it to provide robust estimates of state occupancy probabilities under independent right censoring (Aalen et al. (2001); Datta and Satten (2001)). This property does not hold for the piecewise constant model with interval censored data but for the example there is close agreement between the strictly marginal nonparametric analysis and the model-based analysis for the estimates of the cumulative distribution functions.

Alternative approaches to estimating the baseline intensities are certainly possible including methods based on local likelihood (e.g., Loader (1996); Betensky et al. (1999); Betensky et al. (2002a)). The complete data log likelihood given in Section 2 is amenable to analysis based on polynomial approximations to the intensity functions. The result of this approach would be a further relaxation of the assumptions regarding the baseline functions, and perhaps even closer agreement between the marginal nonparametric estimate and the multistate estimates of the state occupancy probabilities. Although we are currently working on such an implementation, there remains considerable appeal to the methods proposed here because of the straightforward approaches available for variance estimation

and the fact that grid sizes and smoothing parameters (e.g., bandwidth) need not be specified.

6 SUPPLEMENTARY MATERIAL

A SAS macro for fitting the multiplicative model is available under the Paper Information link at the *Biometrics* website <http://www.biometrics.tibs.org>.

ACKNOWLEDGEMENTS

This research was supported by grants from the Canadian Institutes for Health Research and the Natural Sciences and Engineering Research Council of Canada. R.J. Cook is a Canada Research Chair in Statistical Methods for Health Research.

REFERENCES

- Aalen, O. O., Borgan, O., and Fekjaer, H. (2001). Covariate adjustment of event histories estimated with Markov chains: the additive approach. *Biometrics*, 57:993–1001.
- Andersen, P. K., Borgan, O., Gill, R. D., and Keiding, N. (1993). *Statistical Models Based on Counting Processes*. Springer-Verlag, New York.
- Bellamy, S. L., Li, Y., Ryan, L. M., Lipsitz, S., Canner, M., and Wright, R. (2004). Analysis of clustered and interval-censored data from a community-based study in asthma. *Statistics in Medicine*, 23:3607–3621.
- Betensky, R. A. and Finkelstein, D. (2002b). Testing for dependence between failure time and visit compliance with interval-censored data. *Biometrics*, 58:58–63.
- Betensky, R. A. and Finkelstein, D. M. (1999a). An extension of Kendall's coefficient of concordance to bivariate interval censored data. *Statistics in Medicine*, 18:3101–3109.
- Betensky, R. A. and Finkelstein, D. M. (1999b). A non-parametric maximum likelihood estimator for bivariate interval censored data. *Statistics in Medicine*, 18:3089–3100.
- Betensky, R. A., Lindsey, J. C., Ryan, L. M., and Wand, M. P. (1999). Local EM estimation of the hazard function for interval-censored data. *Biometrics*, 55:238–245.
- Betensky, R. A., Lindsey, J. C., Ryan, L. M., and Wand, M. P. (2002a). A local likelihood proportional hazards model for interval censored data. *Statistics in Medicine*, 21:263–275.
- Chen, B. E., Cook, R. J., Lawless, J. F., and Zhan, M. (2005). Statistical methods for multivariate interval-censored recurrent events. *Statistics in Medicine*, 24:671–691.
- Cook, R. J. and Lawless, J. F. (2007). *The Statistical Analysis of Recurrent Events*. Springer, New York.
- Cox, D. R. and Miller, H. D. (1977). *Theory of Stochastic Processes*. Chapman & Hall, London.
- Datta, S. and Satten, G. A. (2001). Validity of the aalen-johansen estimators of stage occupation probabilities and nelson-aalen estimators of integrated transition hazards for non-Markov models. *Statistics and Probability Letters*, 55:403–411.

- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39:1–38.
- Finkelstein, D. (1986). A proportional hazards model for interval-censored failure time data. *Biometrics*, 42:845–854.
- Friedman, M. (1982). Piecewise exponential models for survival data with covariates. *Annals of Statistics*, 10:101–113.
- Gentleman, R. and Vandal, A. (2001). Computational algorithms for censored-data problems using intersection graphs. *Journal of Computational and Graphical Statistics*, 10:403–421.
- Gentleman, R. and Vandal, A. (2002). Nonparametric estimation of the bivariate CDF for arbitrarily censored data. *Canadian Journal of Statistics*, 30:557–571.
- Goggins, W. B. and Finkelstein, D. M. (2000). A proportional hazards model for multivariate interval-censored failure time data. *Biometrics*, 56:940–943.
- Goggins, W. B., Finkelstein, D. M., and Zaslavsky, A. M. (1999). Applying the Cox proportional hazards model when the change time of a binary time-varying covariate is interval censored. *Biometrics*, 55:445–451.
- Grüger, J., Kay, R., and Schumacher, M. (1991). The validity of inferences based on incomplete observations in disease state models. *Biometrics*, 47:595–605.
- Hougaard, P. (2000). *Analysis of Multivariate Survival Data*. Springer, New York.
- Jamshidian, M. and Jennrich, R. (2000). Standard errors for em estimation. *Journal of the Royal Statistical Society, Series B*, 62:257–270.
- Joe, H. (1997). *Multivariate Models and Multivariate Dependence Concepts*. Chapman & Hall, New York.
- Kalbfleisch, J. D. and Lawless, J. F. (1985). The analysis of panel data under a Markov assumption. *Journal of the American Statistical Association*, 80:863–871.
- Kalbfleisch, J. D. and Lawless, J. F. (1989). Some statistical methods for panel life history data. *Proceedings of the Statistics Canada Symposium on Analysis of Data in Time*.
- Kim, M. Y. and Xue, X. (2002). The analysis of multivariate interval-censored survival data. *Statistics in Medicine*, 21:3715–3726.
- Lawless, J. F. and Zhan, M. (1998). Analysis of interval-grouped recurrent-event data using piecewise constant rate functions. *The Canadian Journal of Statistics*, 26:549–565.
- Liang, K. Y. and Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, 73:13–22.
- Lindsey, J. C. and Ryan, L. M. (1993). A three-state multiplicative model for rodent tumorigenicity experiments. *Applied Statistics*, 42:283–300.
- Lindsey, J. C. and Ryan, L. M. (1998). Tutorial in biostatistics: methods for interval-censored data. *Statistics in Medicine*, 17:219–238.
- Loader, C. R. (1996). Local likelihood density estimation. *The Annals of Statistics*, 24:1602–1618.

- Louis, T. (1982). Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44:226–233.
- Meng, X. and Rubin, D. (1991). Using EM to obtain asymptotic variance-covariance matrices: the SEM algorithm. *Journal of the American Statistical Association*, 86:899–909.
- Riggs, B. L., Hodgson, S. F., O’Fallon, W. M., Chao, E. Y., Wahner, H. W., Muhs, J. M., Cedel, S. L., and Melton, L. J. (1990). Effect of fluoride treatment on the fracture rate in postmenopausal women with osteoporosis. *New England Journal of Medicine*, 322:802–809.
- Shih, J. H. and Louis, T. A. (1995). Inferences on the association parameter in copula models for bivariate survival data. *Biometrics*, 51:1384–1399.
- Sun, J. (2006). *The Statistical Analysis of Interval-Censored Failure Time Data*. Springer, New York.
- Turnbull, B. W. (1976). The empirical distribution function with arbitrarily grouped, censored and truncated data. *Journal of the Royal Statistical Society, Series B*, 38:290–295.
- Warkentin, T. E., Cook, R. J., Marder, V. J., Sheppard, J. A., Moore, J. C., Eriksson, B. I., Greinacher, A., and Kelton, J. G. (2005). Anti-platelet factor 4/heparin antibodies in orthopedic surgery patients receiving antithrombotic prophylaxis with fondaparinux or enoxaparin. *Blood*, 12:3791–3796.
- Wei, L. J., Lin, D. Y., and Weissfeld, L. (1989). Regression analysis of multivariate incomplete failure time data by modeling marginal distributions. *Journal of the American Statistical Association*, 84:1065–1073.