STATISTICS IN MEDICINE
Statist. Med. 2001; **20**:2951–2964 (DOI: 10.1002/sim.888)

Event history graphs for censored survival data

Joel A. Dubin[1,*,†], Hans-Georg Müller[2] and Jane-Ling Wang[2]

[1]*Department of Epidemiology and Public Health, Yale University, 60 College Street, New Haven, CT 06520, U.S.A.*
[2]*Department of Statistics, University of California, Davis, Davis, California 95616, U.S.A.*

SUMMARY

A compact graphical device for combining survival and time-varying covariate information is proposed. The proposed graph contains the Kaplan–Meier estimator for right-censored data and a simultaneous display of the behaviour of time-dependent covariate(s) and the lifetime for each subject in the sample. The observed levels of time-dependent covariates are possibly subjected to an initial dimension reduction or smoothing step to produce a continuous covariate function. Values of this function are plotted on a horizontal bar for the length of the lifetime of the subject. Covariate information for censored data is also incorporated. The union of the horizontal bars forms the Kaplan–Meier estimator of the survival function. Our graphical method is implemented with a new S-plus function and demonstrated in several applications. Copyright © 2001 John Wiley & Sons, Ltd.

1. INTRODUCTION

In survival analysis, the Kaplan–Meier estimator provides a non-parametric estimate of the survival function of a specified cohort, naturally incorporating right-censored data. The plot of the Kaplan–Meier estimator is used as a graphical device to demonstrate the survival function of a cohort. The estimator itself has many nice properties such as being the non-parametric maximum likelihood estimator for the true survival function (see, for example, Anderson *et al.* [1]). A graphical limitation of the Kaplan–Meier curve is its inability to include individual-level covariate information. Particularly in the case of time-dependent covariates, such information can be an important part of data exploration and model building.

A recently introduced graphical tool for survival analysis is the event chart (Goldman [2], Dubin *et al.* [3], Lee *et al.* [4]), a graphical technique exhibiting individual-level survival information. With certain implementations of event charts, covariate and survival information can be shown simultaneously. However, event charts are not currently amenable to optimally display time-dependent covariate information.

*Correspondence to: Joel A. Dubin, Department of Epidemiology and Public Health, Yale University, P.O. Box 208034, 60 College Street, New Haven, CT 06520-8034, U.S.A.
†E-mail: joel.dubin@yale.edu

Contract/grant sponsor: National Science Foundation; contract/grant numbers: DMS-9971602, DMS-9803627

Received June 2000
Accepted November 2000

As a new method we propose here the event history graph. It brings together the strengths of both Kaplan–Meier curves and event charts, and simultaneously provides the ability to incorporate time-dependent covariates. A previous version of the event history graph allowed the incorporation of time-dependent covariate information, but did not account for censoring and was restricted to cases where tightly and regularly spaced covariate measurements were available (Carey *et al.* [5]). We propose here a flexible event history graph that incorporates censoring and time-dependent covariate information.

The methodology for the event history graph is described in Section 2. In Section 3, we will provide examples of event history graphs, using data from two separate studies. The first example is the well-known Stanford heart transplantation study that contains right censoring with a binary time-dependent covariate. A second example is a liver cirrhosis study that contains right censoring with a continuous time-dependent covariate. We end with a discussion of various aspects of event history graphs in Section 4. The algorithm used to create the event history graphs and its implementation as an S-plus function is briefly described in an Appendix.

2. THE EVENT HISTORY GRAPH

The event history graph combines information on individual-level time courses with survival information as conveyed by the Kaplan–Meier estimator. The abscissa (x-axis) represents time in the study, while the ordinate (y-axis) represents per cent surviving, as in the customary plot of the Kaplan–Meier estimator. Each individual will have their time course represented as a colour coded (or, alternatively, grey shaded) horizontal bar.

These bars are characterized by length, width (thickness) and colour. All observed times are sorted in ascending order. In case of ties between censored and uncensored observations, the uncensored observations are ordered first. The lengths of all bars are determined by time-to-event (lifetime or censoring time). For censored times, the extension of such a bar to the right is drawn in a 'missing value' colour until the bar hits the vertical jump produced by the next uncensored observation in the Kaplan–Meier estimator. For uncensored observations, the length of the coloured bar corresponds to the survival time. The width (thickness) of the bar depends on the censoring status and the number at risk at the next uncensored observation. The colour of the bars indicates the covariate values and will be discussed momentarily. Note that there is an option to draw a small censoring mark to the right of all bars with censored lifetimes. In some applications, this allows for better visual assessment where censoring occurs (see Figures 4 to 6 and Plates 1 to 3 in Section 3 for examples where censoring is thus marked).

Formally, assume the data are i.i.d. pairs (X_i, δ_i), $i = 1, \ldots, n$, where $X_i = \min(T_i, U_i)$, with $T_i =$ event time and $U_i =$ right-censoring time for subject i. Here, $\delta_i = 1$ if $T_i \leqslant U_i$ and $\delta_i = 0$ if $T_i > U_i$.

Let $T_{(1)}, T_{(2)}, \ldots, T_{(K)}$, $K \leqslant n$, be the ordered observed event times for the n subjects. Also, let n_k be the number at risk for the specified event just before $T_{(k)}$, and d_k be the corresponding number of events at $T_{(k)}$. The Kaplan–Meier estimator [6] for the survival function $\bar{F}$ is then defined as

$$\hat{\bar{F}}(t) = \prod_{k:T_{(k)} \leqslant t} \left(1 - \frac{d_k}{n_k}\right)$$

Now, assume there are l censored observations right before a single uncensored one. Then the Kaplan–Meier jump size is

$$\Delta_{(k)} = \hat{\bar{F}}(T_{(k)}+) - \hat{\bar{F}}(T_{(k)}) = \frac{1}{n_{k+1}}\hat{\bar{F}}(t) \tag{1}$$

at $T_{(k)}$ for an uncensored observation. To determine the bar widths in the event history graph, this jump size is divided into $(l+1)$ equal sized bars, the last bar corresponding to the uncensored observation. That is, each of the bars associated with the kth jump size $\Delta_{(k)}$ has a width of $\Delta_{(k)}/(l+1)$.

We discuss now the incorporation of time-dependent covariates and the utilization of colour in the event history graph. Time-independent covariates are included as a special case and the proposed event history graph will also be of interest for such data. Note that the values of the covariate(s) may be transformed for example, by a logarithmic or Box–Cox transformation, before constructing the event history graph. We are able to easily incorporate a single time-dependent covariate. Sometimes data will include more than one covariate of interest, while a single event history graph is desired. This may be achieved by utilizing dimension reduction techniques. Examples are principal components or single index methods. In fact, any mapping which projects a vector of covariates to dimension one could be considered.

The most convenient dimension reduction scheme for survival data will be to use a Cox [7] proportional hazards regression model. Given p time-varying covariates $Z_1(t),\ldots,Z_p(t)$, the Cox model involves a baseline hazard rate $\lambda_0(t)$ and relates the hazard rate of an individual with covariate functions $Z_1(t),\ldots,Z_p(t)$ to the baseline hazard via $\lambda(t|Z_1,\ldots,Z_p) = \lambda_0(t)\exp(\eta(t))$, where $\eta(t) = \sum_{k=1}^{p}\beta_k Z_k(t)$ (see Andersen *et al.* [1]). This is a single index model with single index $\eta(t)$. In general, a single index is a linear projection used for dimension reduction and corresponds to a linear combination of predictors (as is commonly used for the linear predictor in generalized linear models, see, for example, McCullagh and Nelder [8]).

For multiple covariates, time-varying or not, we then substitute an estimate $\hat{\eta}(t) = \sum_{k=1}^{p}\hat{\beta}_k Z_k(t)$ for $\eta(t)$, obtained by maximizing partial likelihood. This acts then as a one-dimensional time-varying covariate which we may conveniently utilize for the event history graph. For an illustration of this technique, see the example in Section 3.3.

Assuming we have a single time-dependent covariate or have reduced a higher dimensional covariate to dimension one, we will use a colouring scheme, chosen by the user, to indicate the varying levels of the time-dependent covariate over each individual's observed length of time, which is $X_i = \min(T_i, U_i)$. The colour will then signify the magnitude of the covariate at any particular time t, for each individual. One can also implement the method on a grey-scale scheme, where different shades of grey would correspond to the varying levels of the covariate.

Often there exist time gaps in the measurements of the covariate due to uneven designs or more commonly to missing data. A simple remedy is to utilize a smoothing technique, such as local polynomial regression (see, for example, Fan and Gijbels [9]) to obtain a continuous covariate function. Say, for individual i, we have collected m_i observations $y_{i,j}$, $j = 1,\ldots,m_i$ at times $t_{i,j}$ of a time-dependent covariate, Z, over $[0, X_i]$, the total period of observation for the ith subject. Then, using local linear regression, we may obtain $Z(t)$ for a given t,

Table I. Artificial survival data set with ten observations sorted by ascending follow-up time.

Patient	Death during follow-up 1 = yes/0 = no	Follow-up time (in months)	Surgery during follow-up 1 = yes/0 = no	Surgery time (in months)*
1	1	5	0	NA
2	1	7	0	NA
3	0	8	0	NA
4	1	9	1	4
5	1	9	0	NA
6	0	12	1	2
7	1	13	0	NA
8	0	16	1	3
9	1	17	0	NA
10	0	20	1	4

*NA = no surgery during follow-up period.

$0 \leqslant t \leqslant X_i$, by solving the following locally weighted least squares equation:

$$\hat{Z}(t) = \arg\min_a \left(\min_b \sum_{j=1}^{m_i} (y_{i,j} - a - b(t_{i,j} - t))^2 K\left(\frac{t_{i,j} - t}{h}\right) \right) \tag{2}$$

Here, h is a bandwidth to be chosen by the user (automatic selection schemes such as cross-validation are available) and $K \geqslant 0$ is a kernel weight function; an optimal choice is known to be $K(x) = (1 - x^2)1_{[-1,1]}$ (see Müller [10]). We note that there exist fast implementations for local linear fitting [9].

In some cases, the time-dependent covariate may be categorical. For example, covariate values could be 1 = patient underwent surgery prior to time t, 0 = no surgery prior to time t. The covariate may also be a count variable, for example, counting the number of recurrences of a particular cancer during patient follow-up. Another example is an ordered categorical variable, such as clinical staging of cancer. We note that these various types of covariates can be easily incorporated into the methodology for the event history graph.

We have written an S-plus function, *event.history*, to generate event history graphs as described above; the function and corresponding help file may be obtained from the authors. Some details on the algorithm are provided in the Appendix.

3. EXAMPLES AND APPLICATIONS

3.1. Event history graphs in action

Before creating event history graphs from actual study data, we will demonstrate our graphical method with a small artificial survival data set, listed in Table I. This table contains data on ten subjects followed for a maximum period of 20 months. Six subjects experienced the event of interest, that is, death, before the end of the follow-up period, while four subjects are right-censored. There is a single time-dependent covariate, status regarding surgery during the follow-up period. This covariate changes from a value of 0 (no surgery) to 1 at the time

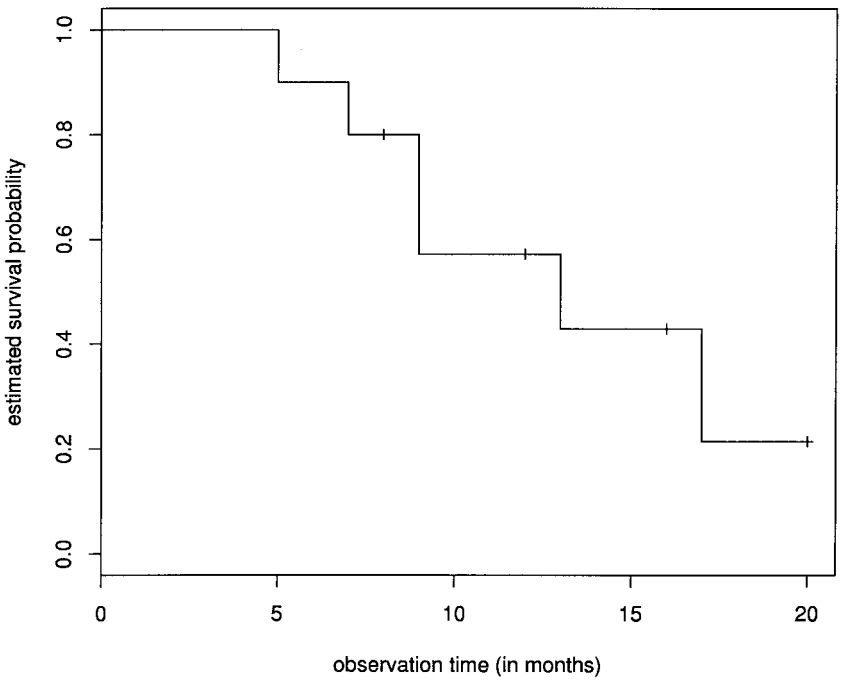

Figure 1. Kaplan–Meier curve for all subjects ($n = 10$) from the artificial data from Table I. Small vertical superimposed dashes along curve represent time of a censored observation.

of surgery. If there is no surgery during follow-up, the covariate remains at 0. Only four of the ten subjects had the surgery.

Figure 1 shows a standard Kaplan–Meier curve for this small cohort. As discussed earlier, the Kaplan–Meier estimator does not convey information regarding a time-dependent covariate. The main feature of the proposed event history graph, shown in Figure 2, is to complement the Kaplan–Meier estimator by providing information both on survival as well as on the time-dependent covariate. For each subject, including subjects with censored lifetimes, the graph contains a bar whose grey scale changes from dark grey to light grey once a subject is treated with surgery.

The hatched bar segments occur only for censored subjects and represent the time from the end of their follow-up period until the next observed event time. The one exception is the last subject, who is censored but has no event time to follow. In this case, a small extension of a censored subject's bar is appended beyond his/her follow-up time, in order to allow recognition that this final bar indeed belongs to a censored individual. The extension of the bar is shown with hatched texture. The Kaplan–Meier curve is then exactly the upper boundary of the union of all bars, discounting the hatched bar extension for the longest observed subject in case the last observation(s) is (are) censored.

This simple example demonstrates the versatility of the proposed event history graphs. These graphs not only convey the usual cohort survival information, but also the time-dependent covariate information for all subjects. By integrating these two features, this graph allows

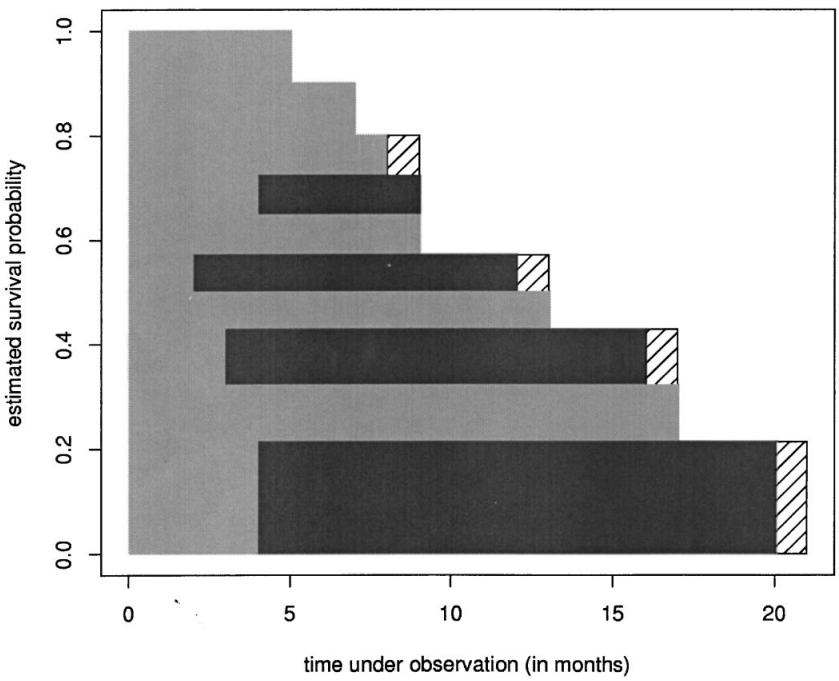

Figure 2. Event history graph for all subjects ($n = 10$) from the artificial data from Table I. Dark grey = no surgery yet during follow-up. Light grey = surgery occurred previously during follow-up. Hatched segments in bars refer to censored subjects.

a visual assessment of whether the time-dependent covariate is related to survival, and it thus complements customary inference procedures such as Cox survival regression with time-varying covariates. Figure 2 provides visual evidence that the subjects who have surgery tend to survive longer. Obviously, if this were a true study, the graph would be complemented by statistical modelling and inference to determine if surgery indeed has a curative effect.

3.2. Application to Stanford heart transplant data

The first real data which we will discuss is from the Stanford Heart Transplant Study; see, for example, Clark *et al.* [11], Crowley and Hu [12] and Kalbfleisch and Prentice [13]. In this study, there was interest in several questions, including determining if a mid-study heart transplant was related to survival. Similar to surgery status in the example from Section 3.1, status regarding heart transplant is the time-dependent covariate here.

Figure 3 displays the Kaplan–Meier curve for the study sample of size $n = 103$, providing information on survival and censoring. The event history graph is seen in Figure 4, and shows the patient transplant history simultaneously with the estimated survival probabilities. It is clear that the vast majority of patients who survived for a long time had received heart transplants early in their follow-up. This is portrayed by the change in bar colour from dark grey (pre-transplant) to light grey (post-transplant). We can also see that the majority of those who died early never received a transplant.

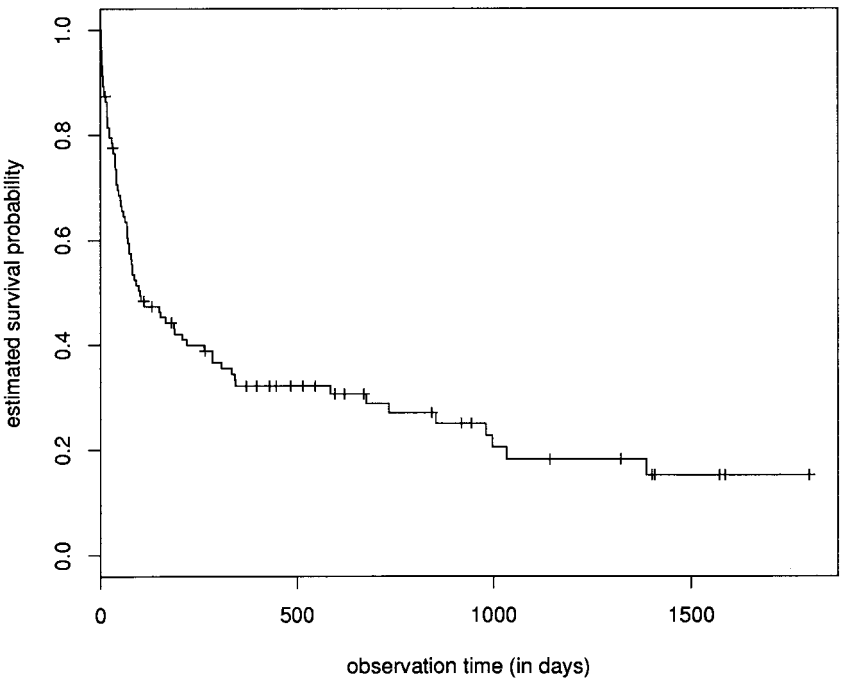

Figure 3. Kaplan–Meier curve for all subjects ($n = 103$) from the Stanford Heart Transplant Study. Small vertical superimposed dashes along curve represent time of a censored observation.

Figures 5(a) and (b) show the event history graphs for the Stanford cohort stratified into two groups by age, <48 and $\geqslant 48$ years old, respectively. Upon visual observation, it appears that the younger patients live longer as compared to the older group. A logrank test comparing these two survival curves yields a borderline statistically significant result with a p-value of $p = 0.09$, with a corresponding 95 per cent confidence interval of $(0.659, 1.040)$ for the parameter $\exp(\beta)$ in the associated Cox model, comparing the younger versus the older group.

Viewing Figures 4, 5(a) and (b) motivates investigating the possibility of an interaction of age and transplant status on survival. In fact the question of this interaction and several other hypotheses have been investigated previously ([12, 13]) using Cox proportional hazards regression and other survival analysis techniques. Our aim here is to show how the event history graph can be used to advantage at the onset of an analysis, exploring the data and possibly generating hypotheses to be tested. In addition, it is useful at the end of an analysis to communicate the results effectively to the medical community.

The Stanford data contained a time-dependent covariate that is relatively easy to handle, that is, a single binary time-varying covariate indicating status of a patient. A more complicated time-dependent covariate would be a continuously measured physiological or clinical variable, which can vary in any direction over time. Such a variable is part of our second study example.

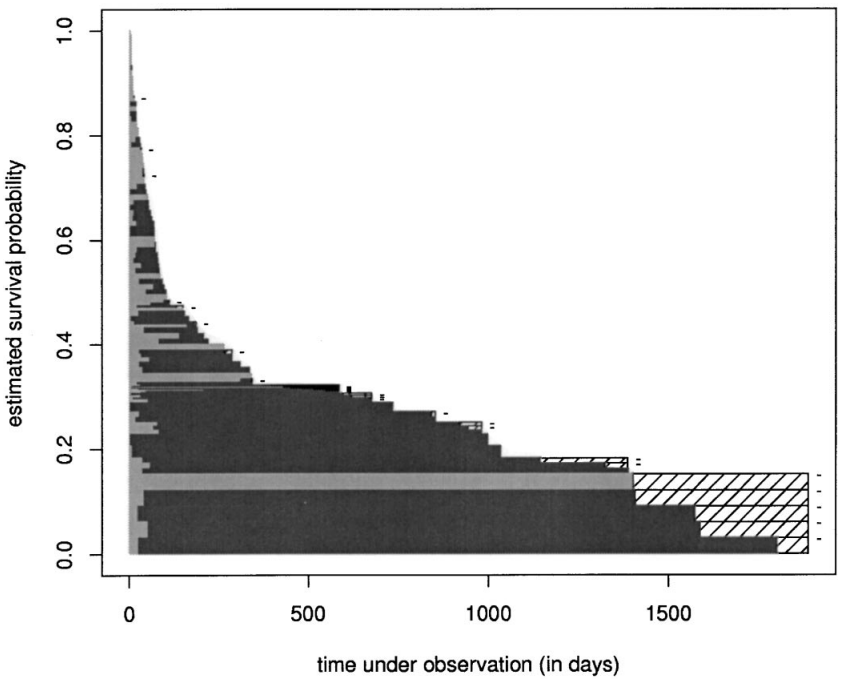

Figure 4. Event history graph for all subjects ($n = 103$) from the Stanford Heart Transplant Study. Dark grey = no transplant yet during follow-up. Light grey = transplant occurred previously during follow-up. Hatched segments in bars, as well as small horizontal dashes following bars, refer to censored subjects.

3.3. Application to liver cirrhosis data

Data from a randomized clinical trial of liver cirrhosis included prothrombin, a circulating blood chemical. This features as a time-dependent continuous covariate, measured between 3 and 17 times per subject over the course of follow-up; the follow-up period ranged from 4 to 4892 days for a total of 488 subjects. The main goal of the study was to see whether the treatment, prednisone, a steroid hormone, was effective in improving survival for cirrhosis patients. This treatment was tested versus a placebo. For further details of the study and formal statistical analyses, see, for example, Andersen *et al.* [1], Schlichting *et al.* [14] and Christensen *et al.* [15].

Event history graphs permit the exploration of the effect of the treatment prednisone and the time-dependent covariate prothrombin on the survival of the cirrhosis patients. Several time-independent variables collected at the beginning of the study were also available for analysis. These included gender, age, histological liver biopsy readings etc.

We constructed five event history graphs for these data. Figure 6 and Plate 1 pertain to the survival and prothrombin levels for all 488 subjects, using different colouring schemes. Plate 2(a) and (b) display survival and prothrombin levels stratified by treatment group. The final graph, Plate 3, provides an example of dimension reduction via the fitting of a Cox proportional hazards regression model, simultaneously incorporating several covariates.

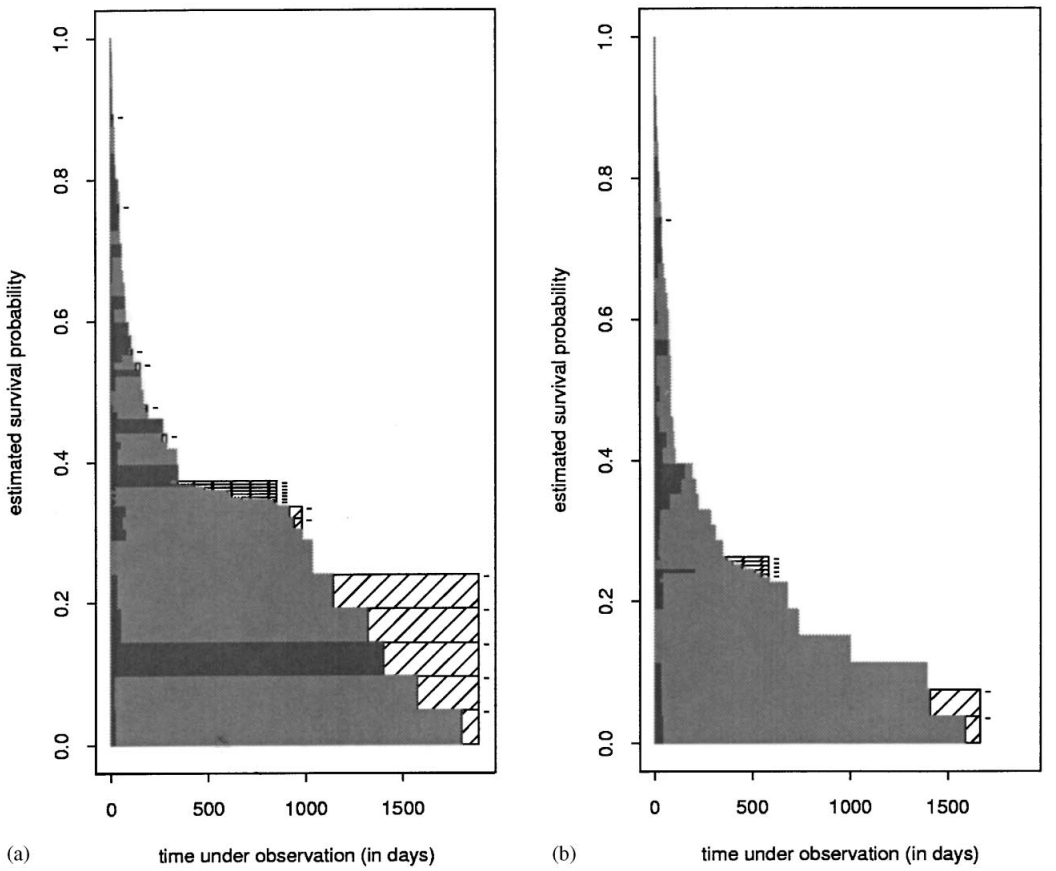

Figure 5. Event history graph (a) for subjects less than 48 years old ($n = 56$) and (b) for subjects greater than or equal to 48 years old ($n = 47$), from the Stanford Heart Transplant Study. Dark grey = no transplant yet during follow-up. Light grey = transplant occurred previously during follow-up. Hatched segments in bars, as well as small horizontal dashes following bars, refer to censored subjects.

The colouring scheme for Figure 6 is in grey scale, with lighter shades of grey corresponding to higher levels of measured prothrombin. Five total levels of prothrombin are indicated, based on intervals defined by quintiles between the observed minimum and maximum prothrombin values in the data set. Prothrombin values were assumed to change at the time of measurement. No smoothing of the prothrombin measurements was performed, though that would be an option for this and similar covariates.

Plate 1 is identical to Figure 6, except that we utilize the colour spectrum instead of shades of grey. Here, royal blue signifies the highest of five levels of prothrombin, while red signifies the lowest. Specifically, the ordering from lowest level to highest, based on quintiles, is red, yellow, green, light blue, royal blue. In general, the user has the option to choose a colouring scheme suitable for the application at hand. Judging from Figures 6 or Plate 1, it appears that those patients surviving longest tend to have higher levels of prothrombin, while, conversely, early death seems to be associated with lower levels of prothrombin.

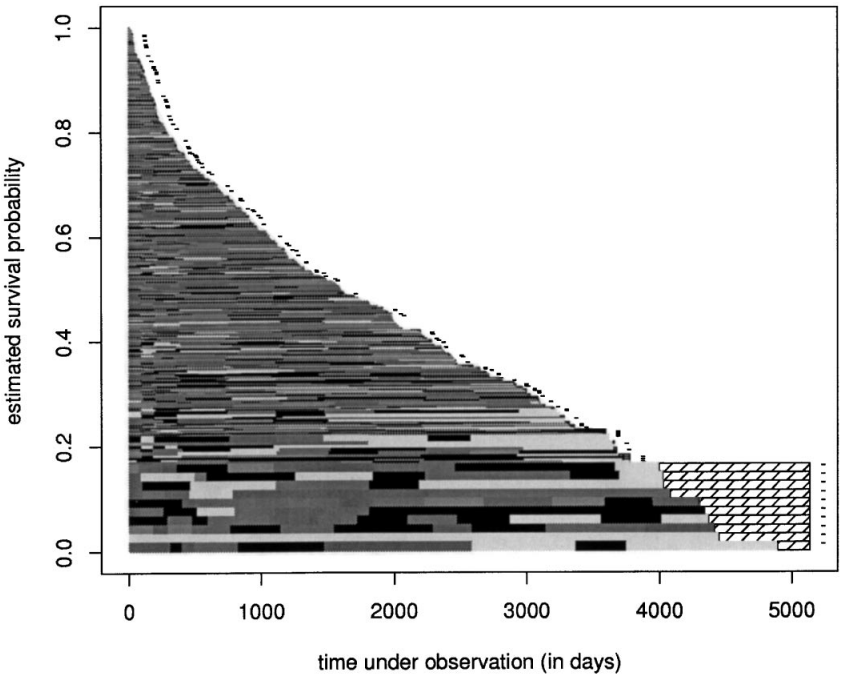

Figure 6. Event history graph in grey scale for all subjects ($n = 488$) from the liver cirrhosis study. Grey scale shading (five total shades) is based on quintiles of prothrombin index. Specifically, darkest grey = lowest prothrombin quintile, darker grey = second lowest quintile, medium grey = middle quintile, lighter grey = second highest quintile and lightest grey = highest quintile. Censoring indicated as in Figure 5.

In Plate 2(a) and (b), we show the event history graphs with prothrombin as covariate, separately for active treatment and placebo, respectively, with colour coding for the subject bars in each graph based on the same overall sample quintile categorization utilized for Plate 1. These graphs point to relatively similar lifetime distributions for the two groups. We note that a logrank test comparing the survival distributions for the two groups is not significant ($p = 0.394$), with an associated 95 per cent confidence interval of $(0.848, 1.070)$ for the parameter $\exp(\beta)$ from the corresponding Cox regression comparing the lifetime of the placebo group to the treatment group.

It does appear that measurements in the lowest quintile of prothrombin (red) are more prevalent for the placebo group, and that measurements in the highest quintile (blue) are more prevalent in the prednisone group. Modelling to determine if an interaction exists between treatment group and prothrombin levels on survival would be useful here. In fact, this question was investigated by Christensen *et al.* [16] and Andersen *et al.* [1]. These researchers determined that an interaction exists between average prothrombin level and treatment. In particular, it was found that prednisone appears to work better for subjects whose average prothrombin levels were in the 70th percentile or higher. Indeed, we can see in Plate 2(a) and (b) that patients who achieve the highest quintile of prothrombin level in the prednisone

Table II. Cox proportional hazards regression for liver cirrhosis data.

Variable	Parameter estimate	Standard error	z-ratio	p-value
Treatment group	−0.1569	0.11882	−1.32	0.1866
Age	0.0591	0.00706	8.37	<0.0001
Prothrombin	−0.0329	0.00225	−14.64	<0.0001
Liver tissue inflammation	−0.4606	0.11848	−3.89	0.0001

group appear to survive at this level for a greater period of time than similar patients in the placebo group.

These graphs illustrate yet another useful feature of event history graphs: they allow us to explore how changes in covariate levels relate to remaining lifetime. A drop in a time-dependent covariate may be an indicator for impending death. There is indeed evidence from the liver cirrhosis study in Plate 2(a) and (b) that a drop in prothrombin levels precedes death for many subjects. Falling prothrombin levels could be a sign of overall failure of body functions, or in itself could be an indicator that cirrhosis is flaring up. Such connections between changing levels and their impact on remaining lifetime are hard to model analytically, and the proposed event history graph allows us to discover such relationships. Emphasis on monitoring the levels of the respective covariate may be a significant clinical consequence of such findings.

Finally, we illustrate dimension reduction with a single index for the case of multiple covariates in Plate 3. The first step is to fit a Cox proportional hazards regression model for time-dependent covariates, as described in Section 2. The following multiple set of covariates was considered: treatment group (0 = prednisone, 1 = placebo); age at entry; prothrombin level; and inflammation of liver tissue (0 = slight or none, 1 = moderate or severe), without consideration of interactions between variables. Only prothrombin level is a time-dependent covariate here; the other predictors are time-independent.

The results from fitting the Cox proportional hazards regression model are in Table II. Higher age, lower prothrombin and greater liver tissue inflammation lead to higher risk of death. Treatment by itself is not statistically significant.

Building on these results from the Cox model, the next step is to generate the single index function, as described in Section 2. Specifically, $\hat{\eta}(t) = \hat{\beta}_1 Z_1 + \hat{\beta}_2 Z_2 + \hat{\beta}_3 Z_3(t) + \hat{\beta}_4 Z_4 = -0.1569(\text{trtgrp}) + 0.0591(\text{age}) - 0.0329(\text{prothrombin}) - 0.4606(\text{liver tissue inflammation})$. Then, for each subject, at each time point t, we calculate $\hat{\eta}(t)$. The resulting univariate time-varying covariate $\hat{\eta}(t)$ is then used as input for the event history graph. This single-event graph is shown in Plate 3, with colour choice again based on the quintiles of the time-varying covariate. Lowest values of the single-index covariate are coded in royal blue and the highest values are coded in red, with the ordering being royal blue, light blue, green, yellow, red. Plate 3 conveys the essence of the Cox model visually. Since three of the four variables in the fitted Cox model have negative coefficients, which outweighs the one variable with a positive coefficient, that is, age, we would expect subjects with low values of $\hat{\eta}(t)$ to live longer. Indeed, Plate 3 reveals a prevalence of royal blue for the longer-lived subjects versus a prevalence of red tones for the short-lived subjects.

4. DISCUSSION AND CONCLUDING REMARKS

We have presented the event history graph, a new graphical method for survival analysis that can display information on survival, censoring and covariates, especially with regard to time-dependent covariates. Through several examples, we have shown that the event history graph is reasonably simple to interpret and can be a valuable exploratory analysis tool for time-to-event studies, in addition to being a compact device for communicating results for time-varying covariates to the medical community.

We note that if in a particular cohort there is a group of patients whose censoring time tends to be greater than the longest observed lifetimes, then the jump sizes at the bottom of the event history graph potentially are much wider vertically than those toward the top. This phenomenon will be magnified when the sample size is large, as is the case for the liver cirrhosis data (see Plate 1). One should then be careful not to inappropriately give disproportionate 'visual' weight to the final observation bars. If this is a problem in certain applications, it is possible to not plot the survival probabilities all the way to zero or to treat the last censored observation as if it were uncensored, as a referee has suggested.

To aid in such a scenario, and also in general to gain additional insights into the data, the user of our *event.history* function can invoke the option of viewing only a specified grid of the event history graph. Hence, one can essentially 'zoom' on a portion of the graph that may otherwise be difficult to visually assess. For example, a user might want to zoom in on a particular quantile of the survival curve. An example of this can be seen in Plate 4, where a zoom of short-term survivors from the liver cirrhosis study is displayed. These are defined by having a value of the Kaplan–Meier estimator of at least 0.9, that is, they constitute the lowest decile in terms of lifetime and are survived by 90 per cent of the subjects in the cohort. This group consists of 73 subjects. Zooming allows to discern finer details of life histories. Note the prevalence of red in Plate 4, confirming that early deaths tend to be associated with high levels of the single index function and low levels of prothrombin. See Section 3.3 for more details.

Another alternative to address the widening of bars towards the bottom would be to consider a graph based on equal width bands. This could help to attenuate the visual influence of wider bands at the bottom of the graph in case of heavy censoring towards the end, as a referee has pointed out. A method for providing equal-width bars already exists, in the form of the event chart ([3, 4]), although time-dependent covariate information is not trivially added to these plots. A major drawback of equal-width bar plots is that the survival function estimate in the form of the Kaplan–Meier curve is lost and therefore the connection between event histories and survival cannot be made anymore. We consider this is to be a central feature of event history graphs.

Combining event history information with the Kaplan–Meier estimator not only allows one to visually assess possible associations of event history with survival, but also provides an information-rich and compact graphical device. Tufte [17] suggests the use of such information-rich graphs and cites as an example a historical graph by Charles Minard which contains six separate items of information in describing the Napoleonean army's march to Moscow. While in some instances the main message of a display may be masked if too much extraneous information is provided, in the case of event history graphs it is not difficult to focus exclusively on the cohort survival information provided in these graphs.

Other colour schemes than the one used in the above examples are possible and will be in some cases preferrable. One could use fewer than the five colours selected in our application

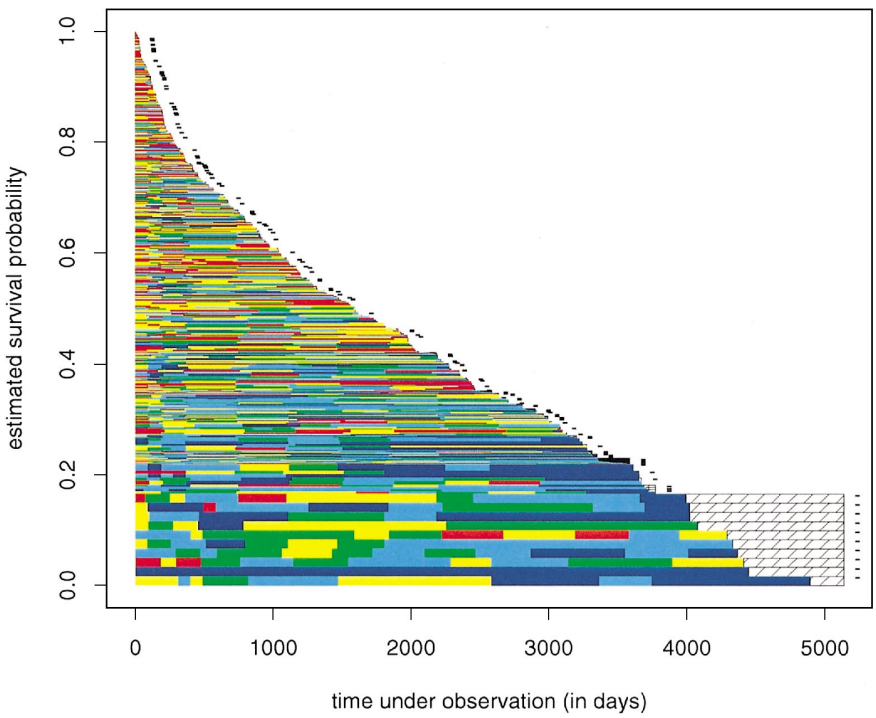

Plate 1. Event history graph in colour for all subjects ($n = 488$) from the liver cirrhosis study. Colour choice (five total colours) is based on quintiles of prothrombin index. Specifically, red = lowest prothrombin quintile, yellow = second lowest quintile, green = middle quintile, light blue = second highest quintile and royal blue = highest quintile. Censoring indicated as in Figure 5.

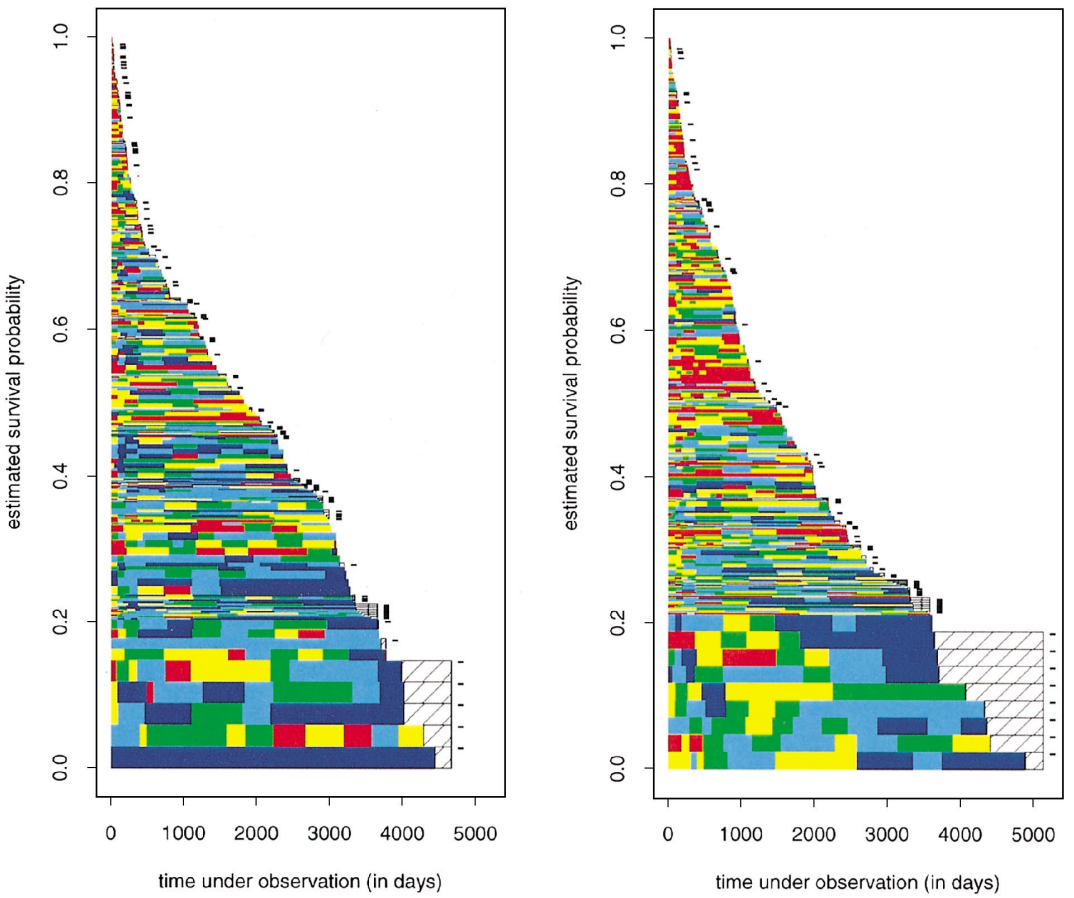

Plate 2. Event history graph (a) for subjects in active treatment (prednisone) arm ($n = 251$) and (b) for subjects in placebo arm ($n = 237$), from the liver cirrhosis study. Colour choice and censoring indicated as in Plate 1.

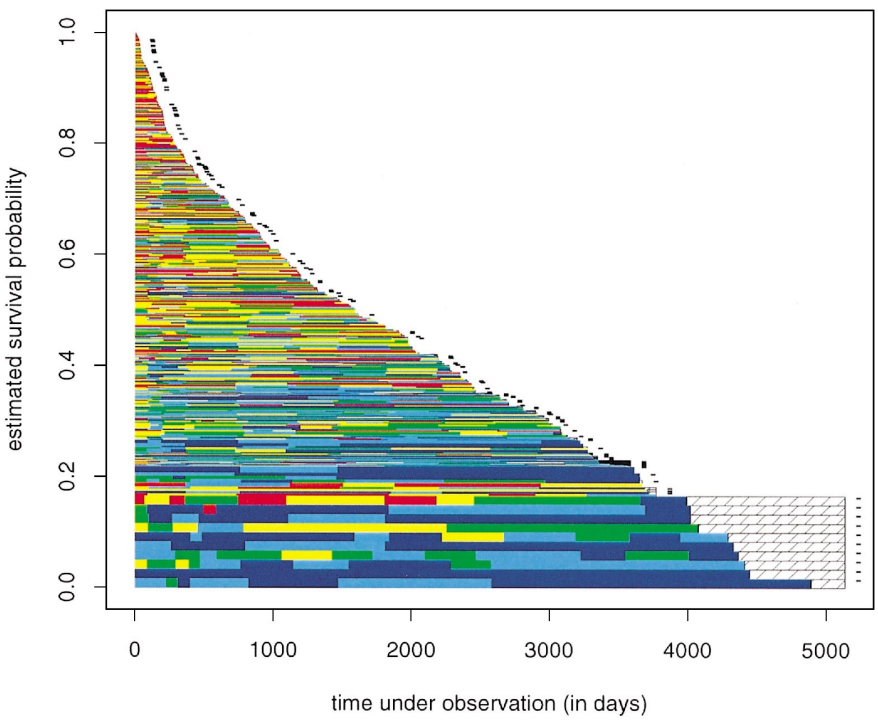

Plate 3. Event history graph for all subjects ($n = 488$) from the liver cirrhosis study based on single index model from Cox proportional hazards regression, using predictor variables listed in Table II. Colour choice (five total colours) is based on quintiles of single index function (that is, $\hat{\eta}$). Specifically, royal blue = lowest single index quintile, light blue = second lowest quintile, green = middle quintile, yellow = second highest quintile and red = highest quintile. Censoring indicated as in Figure 5.

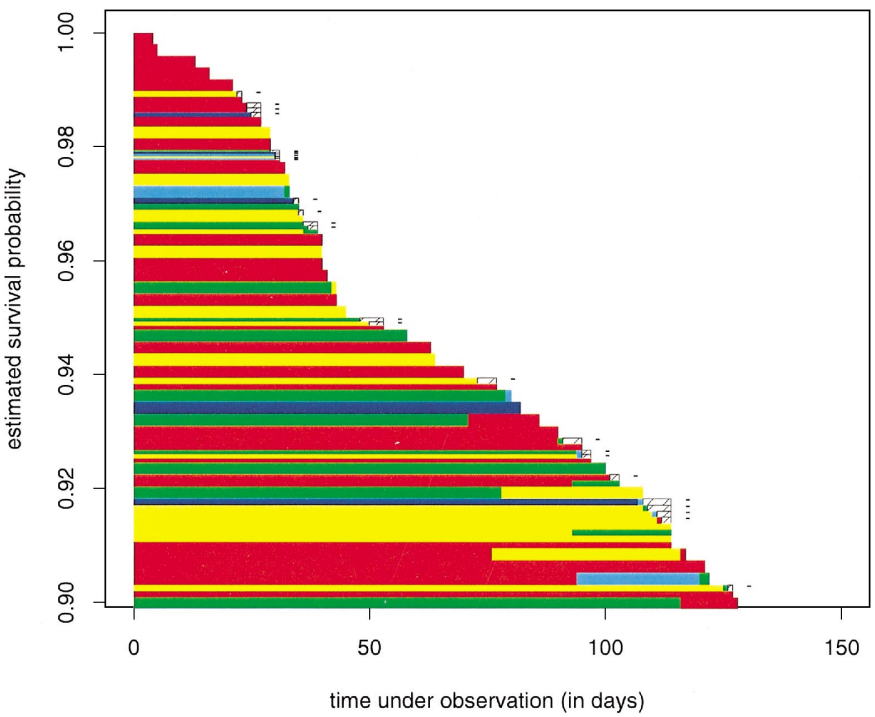

Plate 4. Zoomed version of event history graph from Plate 3, showing the event histories for the subjects in the lowest decile ($n = 73$) for time-to-event.

examples or space the colours closer together in the subjective spectrum, such as varying them between blue and red only. A user can determine the colours of choice in the S-plus function which creates the event history graphs (see Appendix).

In conclusion, the proposed event history graphs serve several functions. They:

1. provide a graphical tool for exploratory data analysis of relations between fixed or time-varying covariates and survival;
2. allow us to explore dynamic changes in a covariate and their relationship with lifetime, such as a plunge in covariate level prior to death;
3. provide a visual assessment of the variability of the covariate time courses between individuals and between cohorts;
4. allow for simultaneous visual comparisons of different cohorts' survival and covariate information;
5. enhance the communication of survival regression inference, such as obtained by applying a Cox proportional hazards regression model, by adding a visual component. This illustrative feature may be of particular interest in regard to medical publications and for communicating results to the medical community.

Especially the last point indicates far reaching potential for the proposed graphical technique. Many publications in the medical literature which contain a survival analysis already include graphs of Kaplan–Meier curves. If desired, these graphs can then be simply extended to include covariate information by using the techniques proposed here. In particular, the event history graph does not require more space than a conventional Kaplan–Meier plot for one cohort, but beyond the Kaplan–Meier curve contains a wealth of additional information.

APPENDIX

We present a summary of the algorithm used to create the event history graph.

1. First, a data set (in S-plus, a data frame) must be created with one row per subject. Within each row should be a subject's observed time, event indicator (for example, 1 = event, 0 = right-censored), and time-dependent covariate values with their associated recorded times.
2. In the program, the data set is ordered by ascending observed time. In case of ties, event times precede censoring times.
3. For uncensored times, count tied event times. For censored times, count the number of other censored times prior to next event time. In case the next event time is tied, also count the number of these tied event times. This is necessary to determine jump sizes and bar heights.
4. Plot bars for each subject, including segments within bars. The lengths of these segments are based on recorded times for covariate changes, and colours of segments are based on the level of the covariate value. The exceptions here are (a) the hatched segments drawn for censored patients from their recorded censoring time to the next uncensored time, (b) if the final unobserved time(s) in the data set is (are) censored, extend a hatched segment by a small amount (for example, 5 per cent) beyond the observed time.
5. In the S-plus program, *event.history*, the user has choices for the following arguments:

   (i) specifying rows of data set to include in graph, say, based on treatment regimen;

(ii) event indicator (for example, 1 = event/0 = censored);
(iii) number of categories for covariate or covariate function;
(iv) how to define covariate or covariate function categories;
(v) colour choices for covariate bar segments – one can choose distinct colours or gradations of grey;
(vi) density of lines in hatched bar segments for censored subjects;
(vii) length of hatched segment extension beyond observed time in case final observation(s) is (are) censored;
(viii) *x*-axis label, *y*-axis label and graph title;
(ix) plot region control, including ability to 'zoom' on specified portion of overall graph.

ACKNOWLEDGEMENTS

We would like to thank P. K. Andersen for providing us access to the liver cirrhosis study data. Also, we would like to thank the two referees for their thoughtful comments, which helped to improve the paper. This research was supported in part by NSF grants DMS-9971602 and DMS-9803627.

REFERENCES

1. Andersen PK, Borgan O, Gill R, Keiding N. *Statistical Models Based on Counting Processes*. Springer-Verlag: New York, 1993.
2. Goldman AI. EVENTCHARTS: Visualizing survival and other timed-events data. *American Statistician* 1992; **46**:13–18.
3. Dubin JA, Lee JJ, Hess KH. The utility of event charts. In *1997 Proceedings of the Biometry Section*. American Statistical Association: Anaheim, CA, 1997; 196–201.
4. Lee JJ, Hess KH, Dubin JA. Extensions and applications of event charts. *American Statistician* 2000; **54**: 63–70.
5. Carey JR, Liedo P, Müller H-G, Wang J-L, Vaupel JW. A simple graphical technique for displaying individual fertility data and cohort survival: case study of 1000 Mediterranean Fruit Fly females. *Functional Ecology* 1998; **12**:356–363.
6. Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* 1958; **53**:457–481.
7. Cox DR. Regression models and life-tables (with discussion). *Journal of the Royal Statistical Society*, *Series B* 1972; **34**:187–220.
8. McCullagh P, Nelder JA. *Generalized Linear Models*. Chapman and Hall: London, 1989.
9. Fan J, Gijbels I. *Local Polynomial Modelling and Its Applications*. Chapman and Hall: London, 1996.
10. Müller H-G. Weighted local regression and kernel methods for nonparametric curve fitting. *Journal of the American Statistical Association* 1987; **82**:231–238.
11. Clark DA, Stinson EB, Griepp RB, Schroeder JS, Shumway NE, Harrison DC. Cardiac transplantation in man, VI. Prognosis of patients selected for cardiac transplantation. *Annals of Internal Medicine* 1971; **75**:15–21.
12. Crowley J, Hu M. Covariance analysis of heart transplant data. *Journal of the American Statistical Association* 1977; **72**:27–36.
13. Kalbfleisch J, Prentice RL. *The Statistical Analysis of Failure Time Data*. Wiley: New York, 1980.
14. Schlichting P, Christensen E, Andersen PK, Fauerholdt L, Juhl E, Poulsen H, Tygstrup N, for The Copenhagen Study Group for Liver Diseases. Prognostic factors in cirrhosis identified by Cox's regression model. *Hepatology* 1985; **3**:889–895.
15. Christensen E, Schlichting P, Anderson PK, Fauerholdt L, Juhl E, Poulsen H, Tygstrup N, for The Copenhagen Study Group for Liver Diseases. A therapeutic index that predicts the individual effects of prednisone in patients with cirrhosis. *Gastroenterology* 1985; **88**:156–165.
16. Christensen E, Schlichting P, Anderson PK, Fauerholdt L, Schou G, Pedersen BV, Juhl E, Poulsen H, Tygstrup N, for The Copenhangen Study Group for Liver Diseases. Updating prognosis and therapeutic effect evaluation in cirrhosis using Cox's multiple regression model for time dependent variables. *Scandinavian Journal of Gastroenterology* 1986; **21**:163–174.
17. Tufte ER. *The Visual Display of Qunatitative Information*. Graphics Press: Chesire, CT, 1983.

