Interactive Dynamic Graphics for Exploratory Survival Analysis

E. Neely Atkinson

The American Statistician, Vol. 49, No. 1 (Feb., 1995), 77-84.

Stable URL:
http://links.jstor.org/sici?sici=0003-1305%28199502%2949%3A1%3C77%3AIDGFES%3E2.0.CO%3B2-8

The American Statistician is currently published by American Statistical Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/astata.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org/
Mon Nov 8 04:33:43 2004

STATISTICAL COMPUTING

This department includes the two sections *New Development in Statistical Computing* and *Statistical Computing Software Reviews*; suitable contents for each of these sections are described under the respective section heading. Articles submitted for the department, outside the two sections, should not be highly technical and should be relevant to the teaching or practice of statistical computing.

Interactive Dynamic Graphics for Exploratory Survival Analysis

E. Neely ATKINSON

The availability of powerful and inexpensive hardware and software for graphical computing makes the use of interactive dynamic graphics feasible for the examination of survival data. The use of these approaches may help the clinical researcher formulate specific hypotheses and select appropriate models. This paper describes some simple techniques, which were implemented in the LISP-STAT environment.

KEY WORDS: Exploratory graphics; Survival data

1. INTRODUCTION

Classic statistical methodology is designed to provide answers to specific questions; exploratory methodology is designed to uncover aspects of data that should generate questions. In collaborative research efforts, the statistician may not be able to determine what these obscure, yet important, aspects are, and may therefore fail to ask questions that could reveal additional information. Thus, it is vital that researchers, who are knowledgeable of the subject matter, but not expert in statistical methodology, have access to technology that allows them to examine data in a revealing, intuitively accessible manner. Interactive dynamic graphics, which enable users to display data and to modify the displays instantaneously, have proven useful for examining different types of data (Cleveland and McGill 1988) and are now included in many commercial statistical packages, e.g., SAS (SAS Institute, Inc. 1990) and S-PLUS (Statistical Sciences, Inc. 1991). Fortunately, clinical researchers generally have access to computer systems that include a high-resolution color monitor and a mouse, features that provide the level of interaction needed to implement interactive dynamic graphics. This functionality, in addition to easily adaptable programming environments, allows data analysts to create sophisticated displays that are appropriate for their goals without relying solely on prepackaged software products.

The purposes of this paper are to (1) demonstrate ways in which dynamic graphics can be used to display survival data and, (2) illustrate that modern hardware and software capabilities enable statisticians who are comfortable with computer systems but not professional programmers to implement such techniques in a reasonable amount of time.

The techniques described in this paper are implemented using the LISP-STAT system. LISP-STAT is a high-level environment for statistical computing available free of charge and compatible with a variety of computer systems (Tierney 1990). All example code for this paper is available from the Statlib server (by ftp to temper.stat.cmu.edu, user name statlib, or by WWW at URL http://lib.stat.cmu.edu/) or by anonymous ftp from odin.mda.uth.tmc.edu.

Table 1. Sample Raw Data for a Hypothetical 60-day Clinical Trial With 12 Patients; Four of the Cases Were Alive at the End of the Trial and Are Censored

Patient	Day Entered	Day of Death	Survival Time
1	19	23	4
2	20	—	>40
3	21	29	8
4	23	38	15
5	25	27	2
6	29	34	5
7	31	59	28
8	34	50	16
9	36	37	1
10	37	—	>23
11	40	—	>20
12	47	—	>13

Table 2. The Kaplan–Meier Life Table Computed From the Example

Survival Time	Number at Risk	Number Dying	Conditional Probability of Death	Conditional Probability of Survival	Unconditional Probability of Survival
1	12	1	.08	.92	.92
2	11	1	.09	.91	.83
4	10	1	.10	.90	.75
5	9	1	.11	.89	.67
8	8	1	.12	.88	.58
13	7	0	.00	1.00	.58
15	6	1	.17	.83	.49
16	5	1	.20	.80	.39
20	4	0	.00	1.00	.39
23	3	0	.00	1.00	.39
28	2	1	.50	.50	.19
40	1	0	.00	1.00	.19

E. Neely Atkinson is Associate Professor, Department of Biomathematics, University of Texas M. D. Anderson Cancer Center, Houston, TX 77030; e-mail: neely@biomath.mda.uth.tmc.edu. This work was supported in part by Grant CA 16672-17 from the National Cancer Institute. The author gratefully acknowledges the helpful comments and suggestions of the reviewers, as well as the substantial programming efforts of Meg Gelder of Rice University.

© 1995 American Statistical Association

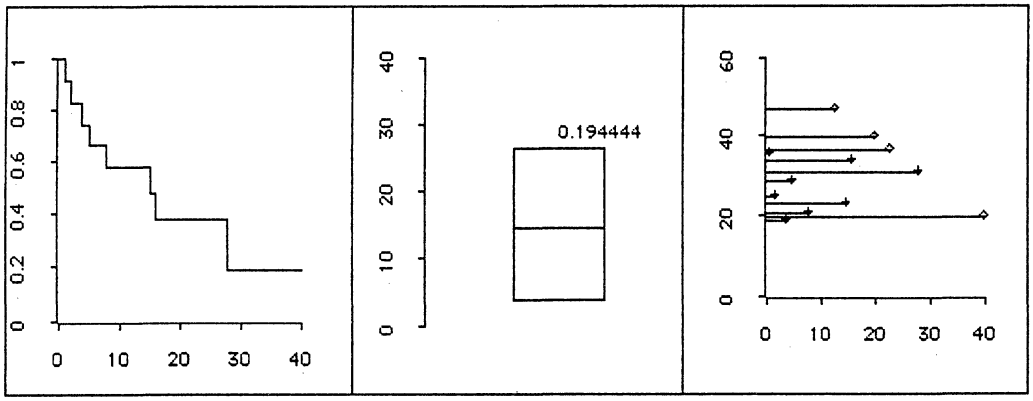

Figure 1. Basic Graphics for Survival Data: The Survival Curve, the Censored Box Plot, and the Eventchart. The survival curve plots time t versus the estimated proportion of subjects surviving until t. The censored box plot shows the times at which 25%, 50%, and 75% of the subjects were alive. The eventchart plots the time at which each subject entered the trial versus the length of time that subject remained on the trial; censored subjects are indicated by a diamond.

2. BASIC GRAPHICS FOR SURVIVAL DATA

The feature that distinguishes survival data from other types of data is the presence of censored values. Suppose we are analyzing time from treatment until death for a group of patients. In most cases, a number of the patients will still be alive at the time of analysis. Consider, for example, the sample data shown in Table 1, which illustrate a hypothetical 60-day clinical trial. Patient 1 entered the trial on day 19 and died on day 23, yielding an observed survival value of 4 days. Patient 2, however, entered the trial on day 20 and was still alive at the end of the trial on day 60. We do not have an observed survival value for this patient, but rather an observed minimum survival value of 40 days. Such data are referred to as censored values. Typical exploratory graphics used to illustrate univariate data (such as histograms or box plots) or multivariate data (such as scattergrams or scatterplot matrices) are not appropriate for illustrating censored data. Because of the presence of censored values, other types of exploratory graphics are needed to display survival data.

The most common way to graphically display survival data is with the survival curve, which is a plot of time (t) versus the estimated proportion of patients surviving until t. The proportion of patients surviving at each time t is usually estimated using the Kaplan and Meier life table method, which uses an argument based on conditional

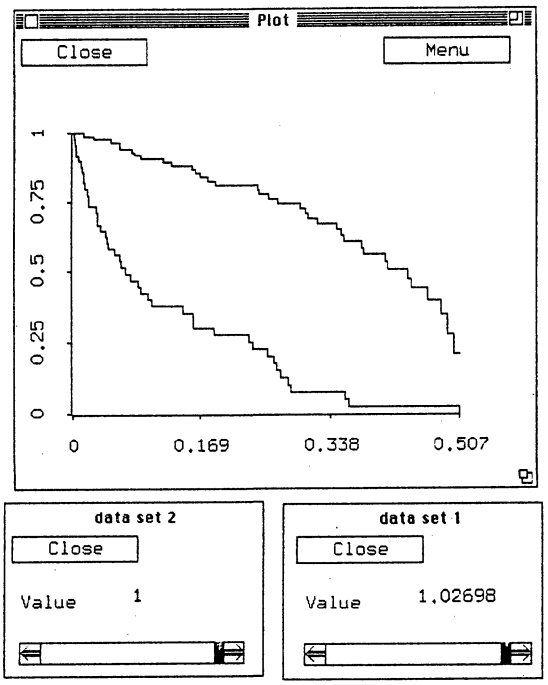

Figure 2. Accelerated Failure Time Curves Showing a Poor Fit. The curves are $S_1(t/w_1)$ and $S_2(t/w_2)$, where $S_1(t)$ and $S_2(t)$ are the Kaplan–Meier estimates of the survival functions for two data sets. The scale factors w_1 and w_2, labeled "value" on the sliders, are adjusted by using the sliders. If the relation between the data sets can be accurately described using an accelerated failure time model, then values for w_1 and w_2 can be found which bring the survival curves into approximate alignment.

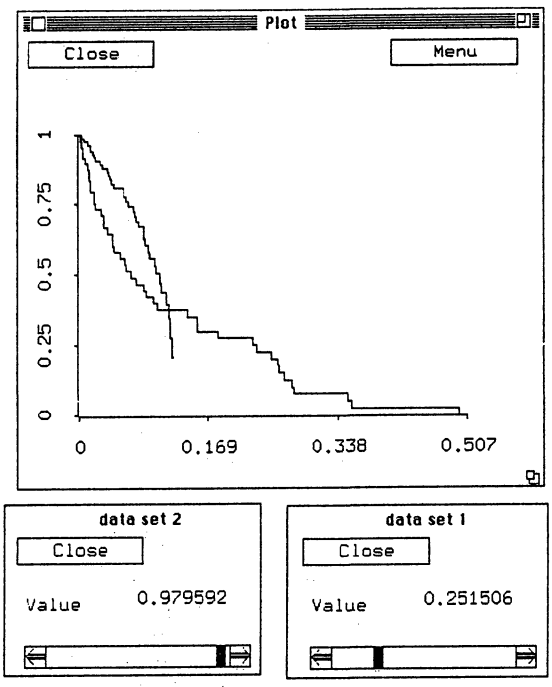

Figure 3. Accelerated Failure Time Curves Showing an Improved Fit. The data sets from Figure 2 are shown with values for w_1 and w_2 adjusted to bring the curves into closer alignment. An accelerated failure time model is probably not appropriate for these data.

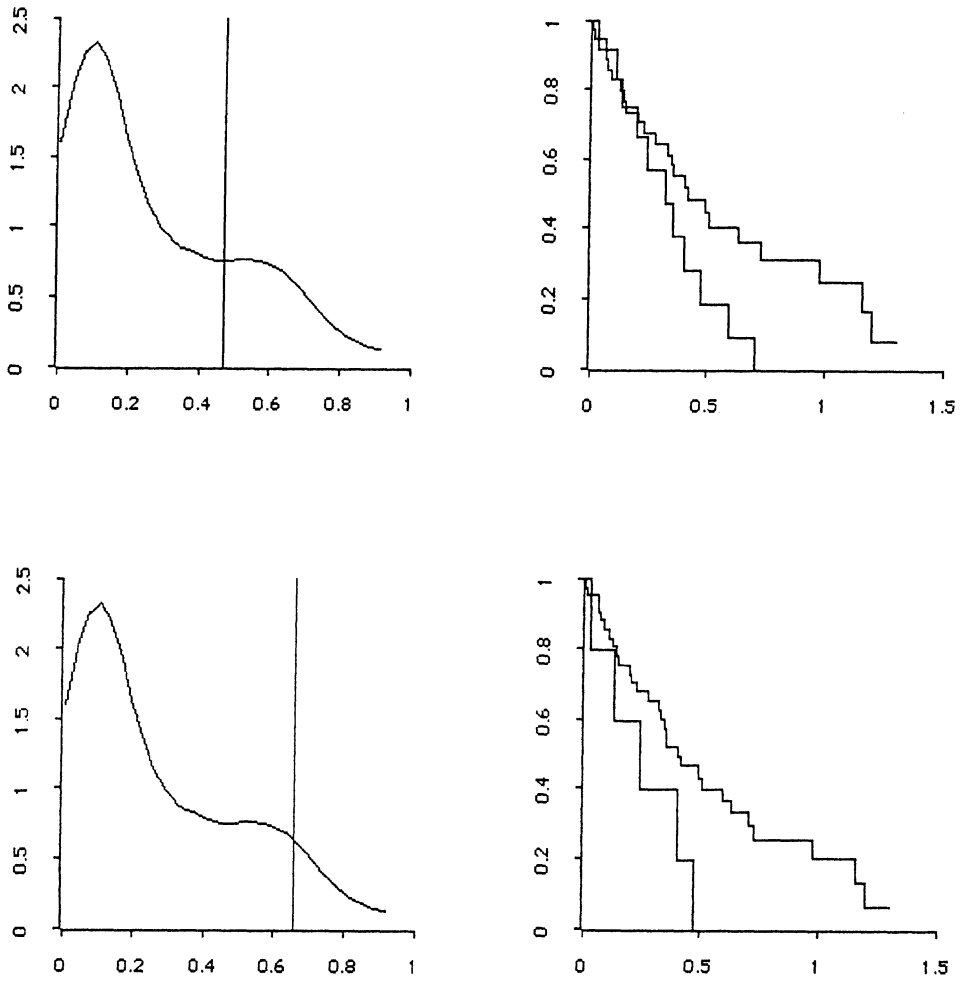

Figure 4. Dichotomizing a Continuous Covariate. This figure shows two displays (upper and lower) illustrating a method for investigating the effects of dichotomizing a continuous covariate. Each display consists of a density estimate for a continuous covariate, with a cut point marked by a vertical line, and a pair of survival curves representing the cases with values of the covariate above and below the chosen cut point. The cut point is adjusted by use of a slider, which is not shown in the figure. In the upper display, the cut point is just under .5. In the lower portion, the cut point has been moved to just over .6, providing better separation of the curves. By moving the slider, the investigator can examine the effects of varying the cut point, thus studying how well the cut point is determined and how much information is being lost by dichotomizing the covariate.

probabilities to include information provided by the censored cases (Kaplan and Meier 1958). A life table computation of the data given in Table 1 is illustrated in Table 2 (p. 77). A survival curve for this life table is shown in the leftmost panel of Figure 1.

Another way to display survival data is with the censored box plot (Gentleman and Crowley 1991). Censored box plots are similar to regular box plots, except that the censored box plot is constructed from a life table, and the regular box plot is constructed from raw data. Horizontal bars are placed at the median and the upper and lower quartiles, if these exist. The value of the survival curve is printed at the point corresponding to the last uncensored time. If a certain quartile has not yet been reached, the sides of the bar are extended to the last uncensored time. Censored box plots contain less information than survival curves; however, with the censored box plot, you can quickly and easily compare groups by visual examination. The censored box plot for the life table in Table 2 is shown in the middle panel of Figure 1. We can see from this box plot that: (1) the median survival is about 15 days, (2) about 75% of the patients survived at least 4 days, (3) about 25% of the patients survived at least 25 days, and (4) the last death occurred at about 28 days at which point the value of the survival curve was .19.

A third way to display survival data is with the eventchart (Goldman 1992). The eventchart contains more information than the survival curve, and can reveal features of the data that are hidden on survival curves and censored box plots; however, the eventchart obscures other features of the data, and in this way complements survival curves and censored boxplots. On an eventchart, the vertical axis represents when the patient is entered onto the study; the horizontal axis represents the length of time the patient is alive and on the study. A horizontal line joins the patient's entry onto the study to the patient's time of death or last follow-up. Different symbols or line types can be used to distinguish between censored and uncensored cases. In the construction of the life table, all patients are considered to have entered the study simultaneously, that is, it is assumed that survival experience does not change over the life of the study. This assumption is not made in the construction of the eventchart. An eventchart for the life table in Table 2 is shown in the rightmost panel of Figure 1.

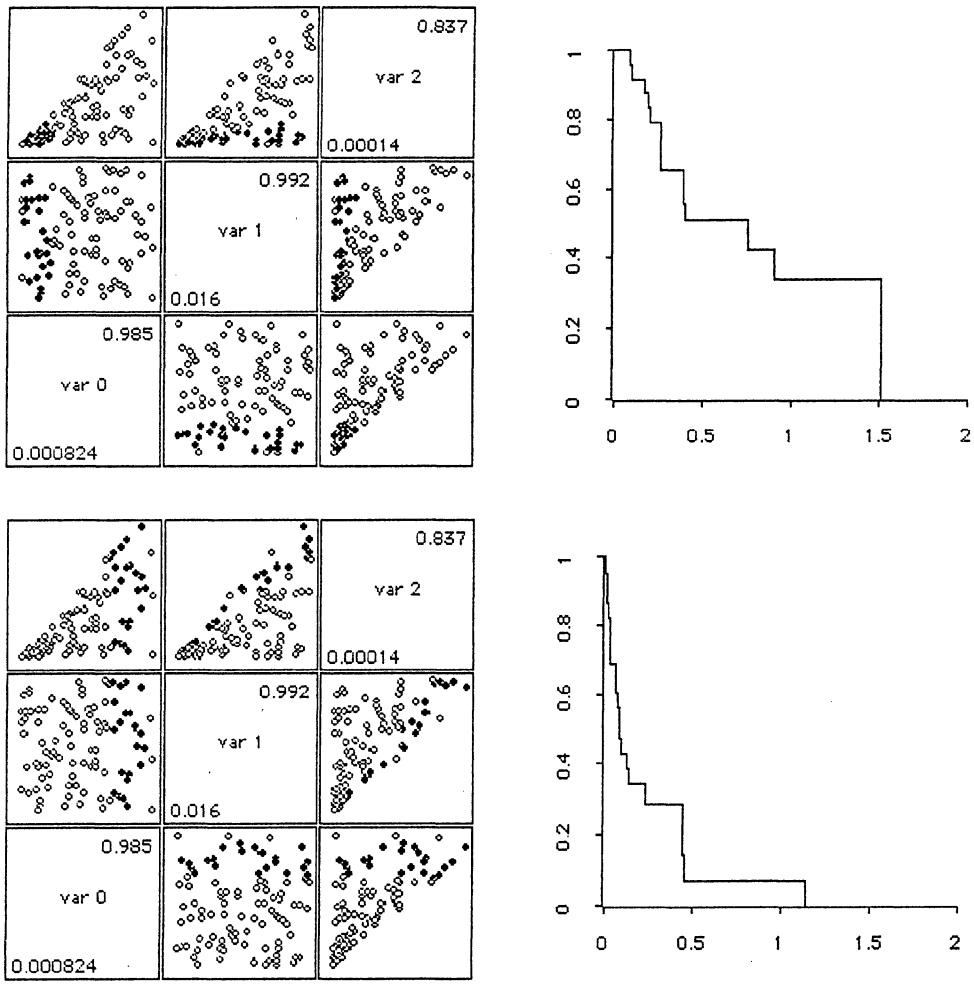

Figure 5. Scatterplot Matrices Linked to Survival Curves. This figure shows two displays (upper and lower) illustrating scatterplot matrices linked to survival curves. The scatterplot matrix in each display shows all pairwise scatterplots for a set of covariates. A subset of cases may be selected from any scatterplot using the mouse, and the selected cases will be highlighted in all the scatterplots. The accompanying survival curve is computed from the selected cases. In the upper display, cases with low values of variable 0 have been selected; in the lower display, cases with high values. We observe that survival decreases with increasing values of var 0. We also observe, by examining the plots of var 1 versus var 2, that the relation between var 1 and var 2 depends on the value of var 0.

Censored points are indicated by open diamonds and uncensored points are indicated by crosses.

3. SOME GRAPHICS FOR REGRESSION MODELS

Two regression models are commonly used to analyze survival data: the proportional hazards model (Cox 1972) the accelerated failure time model (Kalbfleisch and Prentice 1980). The proportional hazards model postulates that $h(t \mid \mathbf{x} = h_0(t)e^{\beta' \mathbf{x}}$, where $\mathbf{x}'$ denotes the transpose of $\mathbf{x}$, $h(t \mid \mathbf{x})$, is the hazard rate at time t given covariate values $\mathbf{x}$, $h_0(t)$ is the underlying hazard rate when all covariates are 0, and β is a set of regression coefficients to be estimated. The accelerated failure time model assumes that $\ln(T) = \mathbf{x}'\beta + \ln(T_0)$, where T is a random variable representing the survival time of patients with covariate values $\mathbf{x}$, T_0 is the underlying survival for patients with covariates values of 0, and β is a set of regression coefficients. In mathematical terms, these models generally carry little intuitive appeal for clinicians; however, they can be given a simple graphical interpretation. Let $S(t \mid \mathbf{x})$ denote the survival function at time t for patients with covariate values $\mathbf{x}$, and let $S_0(t)$ denote the survival function at time t for patients with covariate values of 0. Then, according to the proportional hazards model, $S(t \mid \mathbf{x}) = S_0(t)^{\exp(\beta' \mathbf{x})}$, and according to the accelerated failure time model, $S(t \mid \mathbf{x}) = S_0(e^{\mathbf{x}'\beta}t)$. Thus, the proportional hazards model assumes that the survival curve for a given patient is equal to a basic survival curve raised to a power determined by the patient's covariate values, while the accelerated failure time model assumes that all patients traverse the same survival curve but at time scales adjusted according to each patient's covariate values. Consider the case in which $\mathbf{x}$ is a single binary value, so there are only two curves. If the proportional hazards model assumption holds, the curves can be moved up and down (in a suitably constrained fashion) until they are approximately aligned; if the accelerated failure time model assumption is appropriate, then the curves can be moved left and right (in a constrained fashion) until they are approximately aligned.

Figures 2 and 3 (p. 78) show dynamic plots illustrating this point for the accelerated failure time model. Each plot consists of three components: a graph showing two (scaled) survival curves, and two sliders. The survival curves are $S_1(t/w_1)$ and $S_2(t/w_2)$, where $S_1(t)$ and $S_2(t)$

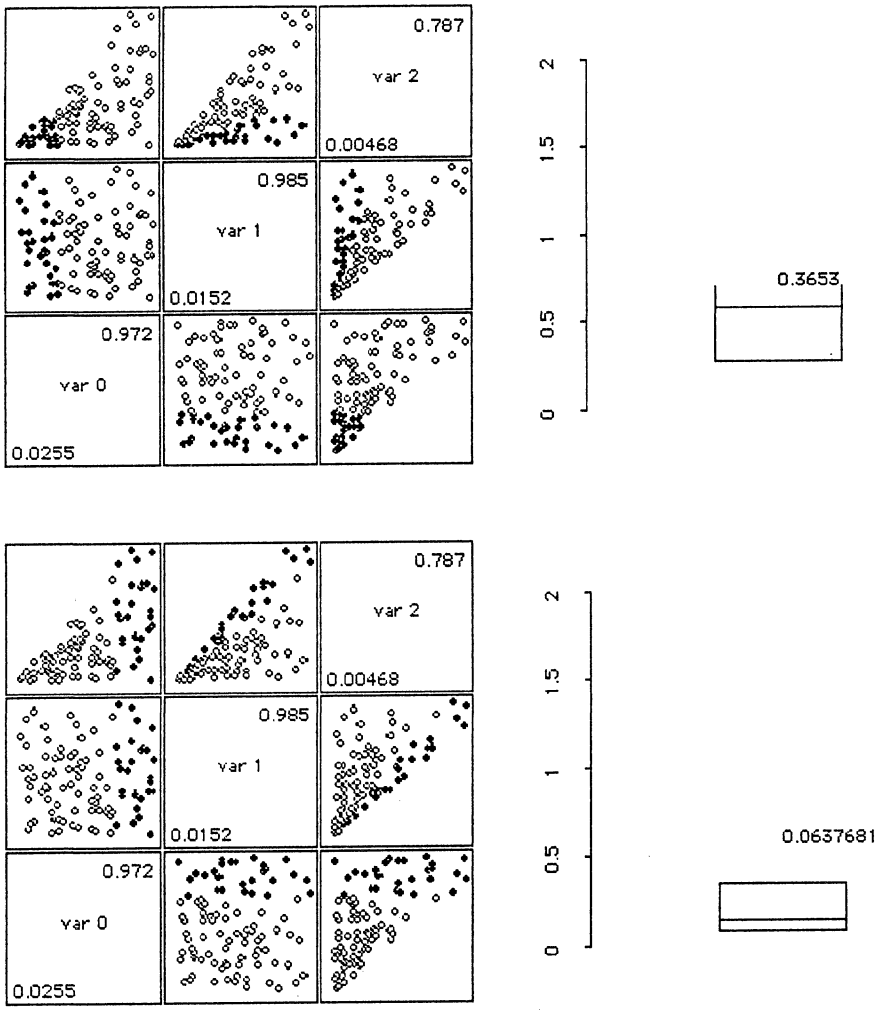

Figure 6. Scatterplot Matrices Linked to Censored Box Plots. This figure shows two displays (upper and lower) illustrating scatterplot matrices linked to censored box plots. In the upper display, cases with low values of var 0 have been selected in a scatterplot matrix and the corresponding censored box plot constructed and displayed; the median survival time for these cases is over .5. In the lower display, cases with high values of var 0 have been selected and the median survival has dropped to under .2.

are the Kaplan–Meier estimates of the survival functions for the two data sets. The values of w_1 and w_2 are varied by adjusting the sliders. In Figure 2, w_1 and w_2 are set to their default values: w_2 is 1 and w_1 is set so that the two survival curves span the same amount of time. In Figure 3, w_1 has been decreased so that the uppermost curve is shifted to the left, bringing the curves closer to alignment. An accelerated failure time model is probably not appropriate for fitting these data.

Although an actual analysis will almost always contain several covariates, none of which need be binary, these plots can still be of use. Assume that the regression model of interest has been fit, yielding estimated coefficients $\mathbf{b}$. Then patients with similar values for $\mathbf{x}'\mathbf{b}$ can be grouped, producing a small number of curves that may then be examined in the fashion described; such a grouping has been used, for example, by Tanner and Wong (1983). The plot could be easily extended to include more than two curves, if this were desired.

Several additional extensions would also be straightforward to implement. First, these plots could also easily be combined to permit the curves to be moved either up and down or back and forth, making it easier to select between the two models. Second, some measure of agreement between the two curves could be added to the display, perhaps a Kolmogorov–Smirnov or log rank statistic. The value of the statistic could be displayed either in a dialog box or graphically as, for example, the height of a bar. It would, of course, be simple to display the p value for the selected statistic, but this raises difficult questions of interpretation. This issue is discussed in more detail in the following section.

4. DICHOTOMIZING A COVARIATE

A common request from clinicians to data analysts is to assist in determining which values of a predictor are "normal" and which are "elevated." A dynamic graph, which assists in this task, is illustrated in Figure 4 (p. 79). This graph consists of two parts: (1) a density plot, which shows a nonparametric density estimate of the predictor with the current cutoff value marked by a vertical line, and (2) a pair of survival curves representing cases with values above and below the selected cutoff point. By adjusting a slider (which is not shown in the figure), the user attempts to find values that separate the cases into those with good and bad prognoses. In the upper pair of plots, the cutoff point is at the midpoint between the upper and

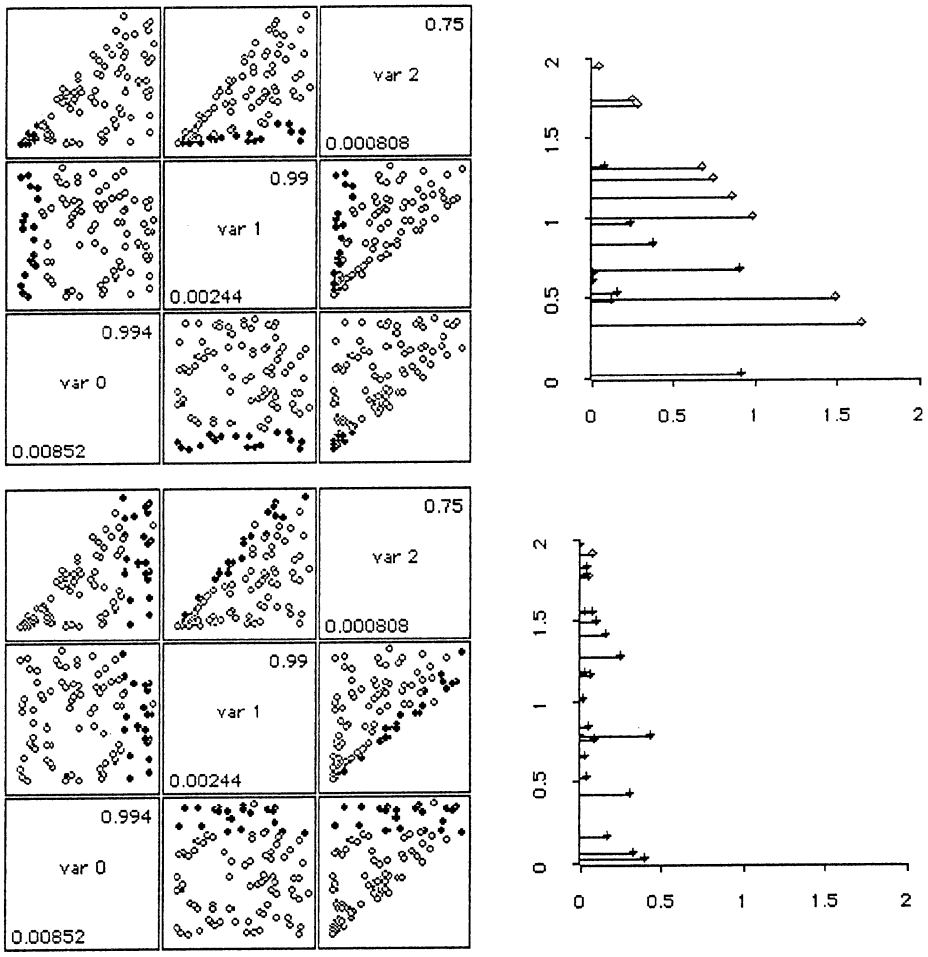

Figure 7. Scatterplot Matrices Linked to Eventcharts. This figure shows two displays (upper and lower) illustrating scatterplot matrices linked to eventcharts. In the upper display, cases with low values for var 0 have been selected in a scatterplot matrix and the corresponding eventchart constructed and displayed. In the lower display, cases with high values of var 0 have been selected. Although it is difficult to estimate survival at a given time from an eventchart, it is clear that cases with low values of var 0 have longer survival times. There is no sign of any difference in the arrival times of cases with low or high values of var 0; this assumption is impossible to check from survival curves or censored box plots.

lower extremes of the covariate; in the lower plots, the cutoff has been moved to the right, giving a somewhat better separation between the curves.

This graphic is useful in selecting a cutoff point for dichotomizing a predictor; more importantly, it is useful for exploring the effects of varying that choice and, indeed, for examining whether such a simplification of the predictor is a good idea at all. If the clinician can see that prognosis varies continuously in relation to the value of the predictor, then it becomes clear that significant and useful information is being discarded by dichotomizing the covariate. Furthermore, the use of this technique indicates how sharply determined the cutoff point is and what the effects of varying the cutoff are, giving the clinician an appropriate sense of either confidence or skepticism.

One difficulty with this technique is that the visual comparison of the two curves is sometimes misleading. When there are few cases in one group, the curves may appear very different, but the variability of the curve associated with the smaller group may also be very high. One way around this would be to display confidence limits for the curves; this is discussed in Section 6. Another possibility would be to display a measure of the agreement between the curves, such as a Kolmogorov–Smirnov or a log rank statistic. Such a display would be useful, but would require extreme caution in its interpretation. In particular, the usual p values associated with these statistics would not be appropriate since varying the cut point is, in effect, performing repeated hypothesis testing (Altman, Lausen, Sauerbrei, and Schumacher 1994).

5. EXAMINING MULTIPLE COVARIATES

Typically, an analysis will examine the effects of several covariates simultaneously. The scatterplot matrix, which shows all pairwise scatterplots for the covariates, is particularly useful for aiding in such an analysis. The scatterplot is extended to a dynamic interactive form by allowing the user to select points from any of the plots, either one by one ("selecting") or in a continuously varying fashion ("brushing"). The selected points are highlighted not only in the panel chosen, but in each of the scatterplots, permitting the relationships among several sets of variables to be examined simultaneously (Cleveland and McGill 1988, pp. 13–22). To examine the effect several covariates have on survival, the scatterplot matrix can be linked to the appropriate survival data display: survival curve, censored box plot, or eventchart. As the selected points change, the survival display can be continuously updated. Using this technique, it is possible to search for subsets of cases that behave in a surprising fashion, which may suggest further

```
;
; performs Kaplan-Meier estimation.
;
(defun kmest (time status)
  (let* ((n (length time))
         (s 1)
         (surv nil))
    (dotimes (i n)
             (setf s (* s (- 1 (/ (select status i) (- n i)))))
             (setf surv (cons s surv)))
    (reverse surv)))
;
; transforms results of kmest into x,y
; coords for plotting.
;
(defun km-plot (time status surv)
  (let* (
         (n (length time))
         (rep-pat (repeat 2 n))
         (time-plot (cons 0 (repeat time rep-pat)))
         (surv-plot (reverse
                     (cons (select surv (- n 1))
                           (repeat (reverse (cons 1
                                                  (select
                                                   surv
                                                   (iseq (- n 1)))))
                                   rep-pat)))))
    (list time-plot surv-plot)))
```

Figure 8. Code for a Scatterplot Matrix Linked to a Survival Curve.

investigations. Figure 5 (p. 80) illustrates scatterplot matrices linked to survival curves. In the upper pair of plots, low values of variable 0 have been selected; in the lower plots, high values have been selected. It is clear from the associated survival curves that higher values of variable 0 are associated with decreased survival. By examining the scatterplot matrix, it is also clear that the relation between variable 1 and variable 2 depends on the value of variable 0.

Figures 6 and 7 (pp. 81 and 82, respectively) show censored box plots and eventcharts (respectively) that are similarly linked to scatterplot matrices.

6. DIFFICULTIES AND EXTENSIONS

All of the displays discussed in this article share a common flaw: they do not include information on the variability of the estimated survival. Consider, for example, the use of a slider to choose a cutoff value for a covariate. Since the survival curve is better determined at the beginning, where there are more cases, separations near the beginning of the curve should carry more weight. Furthermore, when the selected cutoff value is set near one of the covariate extremes, one of the curves will be based on a small number of cases, and may be therefore unreliable. One way to avoid these difficulties would be to add upper and lower confidence intervals to the displayed curves. Those additional data, however, produce a busy display that may be difficult to interpret. This difficulty might be reduced by using shaded regions rather than simply adding additional curves to the graphs. Another possibility would be to provide pointwise confidence intervals for any point on which the user clicks. In addition to the problems caused by overly busy graphs, it is our experience that the user's eye is drawn to the position on the curve where the confidence intervals are widest, which in fact are the regions that should receive the least attention. Perhaps other researchers will be successful in devising methods to effectively display variability on these graphs.

```
;
; generates a scatterplot matrix linked to
; a Kaplan-Meier estimate.
;
(defun scat-km (time status &rest covars)
  (let* ((scat (scatterplot-matrix covars))
         (data (km-plot time status (kmest time status)))
         (km (plot-lines (first data) (second data) :title "KM Plot"))
         (current-points (iseq (length time))))
    (send km :range 1 0 1)
    (send scat :point-state (iseq (length time)) 'hilited)
    (defmeth scat :unselect-all-points ()
      (setf current-points nil)
      (send km :clear)
      (call-next-method))
    (defmeth scat :adjust-points-in-rect (x1 y1 w h s)
      (let* ((p-i-r (send self :points-in-rect x1 y1 w h)))
        (setf current-points
              (case (send self :mouse-mode)
                (brushing p-i-r)
                (selecting (union p-i-r current-points))))
        (if current-points
            (prog*
             ((points (sort-data current-points))
              (time-sel (select time points))
              (status-sel (select status points))
              (new-plot (km-plot time-sel status-sel
                                 (kmest time-sel status-sel))))
             (send km :clear :draw nil)
             (send km :add-lines
                   (first new-plot) (second new-plot))))
        (call-next-method x1 y1 w h s)))
    scat))
```

Figure 9. Code for a Scatterplot Matrix Linked to a Survival Curve (Continued).

7. PROGRAMMING NEW DISPLAYS

The primary purpose of this article is not to announce a complete package for the exploratory analysis of survival data, but rather to demonstrate that such displays are useful and not difficult to produce. Figures 8 and 9 show the complete LISP-STAT code required to generate a scatterplot matrix linked to a survival curve. Although the code may seem opaque to readers who are unfamiliar with LISP-STAT, it should be evident that it is possible to produce sophisticated dynamic graphics with a program of reasonable length. It should also be noted that several of the routines are generally useful and can be shared with other graphics. As a user builds up a library of code, new graphs become simpler. Finally, LISP-STAT contains a powerful set of graphics techniques as primitives of the language. Thus, for example, given a scatterplot matrix linked to a survival curve, no additional programming is required to add a spinning three-dimensional plot also linked to the scatterplot since the three-dimensional plot and the link to the scatterplot are already provided.

8. CONCLUSION

With modern computing environments, it is now possible for researchers who are not expert computer scientists to produce useful and effective dynamic graphics for the display of survival data with a reasonable amount of effort. The availability of such methods should permit the clinical researcher to become a more active participant in the search for aspects of the data that are worthy of further study.

As more researchers become aware of the potential of systems such as LISP-STAT, more sophisticated and informative graphical techniques are sure to be developed. Ideally, the researchers who develop such methods will share them with the research community at large, preferably through such electronic media as Statlib. This paper is intended as an initial effort in that direction.

[*Received November 1993. Revised August 1994.*]

REFERENCES

Altman, D. G., Lausen, B., Sauerbrei, W., and Schumacher, M. (1994), "Dangers of Using 'Optimal' Cutpoints in the Evaluation of Prognostic Factors," *Journal of the National Cancer Institute*, 86, 829–835.

Cleveland, W. S., and McGill, M. E. (ed.) (1988), *Dynamic Graphics for Statistics*, Belmont CA: Wadsworth and Brooks/Cole.

Cox, D. R. (1972), "Regression Models and Life Tables" (with discussion), *Journal of the Royal Statistical Society*, Ser. B, 34, 187–220.

Gentleman, R., and Crowley, J. (1991), "Graphical Methods for Censored Data," *Journal of the American Statistical Association*, 86, 678–683.

Goldman, A. (1992), "EVENTCHARTS: Visualizing Survival and Other Timed-Events Data," *The American Statistician*, 46, 13–18.

Kalbfleisch, J. D., and Prentice, R. L. (1980), *The Statistical Analysis of Failure Time Data*, New York: John Wiley.

Kaplan, E. L., and Meier, P. (1958), "Nonparametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, 53, 457–481.

SAS Institute, Inc., *SAS/INSIGHT User's Guide* (1990), Carey, NC: Author.

Statistical Sciences, Inc., *S-PLUS User's Manual*, Vol. 1, Version 3.0 (1991), Seattle, WA: Author.

Tanner, M. A., and Wong, H. W. (1983), "Comments on 'A Reanalysis of the Stanford Heart Transplant Data,'" *Journal of the American Statistical Association*, 78, 286–287.

Tierney, L. (1990), *LISP-STAT: An Object-Oriented Environment for Statistical Computing and Dynamic Graphics*, New York: John Wiley.

Statistical Computing Software Reviews

RICHARD GOLDSTEIN, Section Editor

This section is similar in organization to a book-review section in other journals; however, software of interest to statisticians is the subject of review here. Emphasis is on software for microcomputers. Programs that operate only in larger mainframe computers will seldom receive review. Normally, producers of programs make a copy of their product available to the section editor, who then selects one or more persons to test the product and prepare a review.

Producers of computer software who wish to have their product reviewed are invited to contact Section Editor Richard Goldstein, Qualitas, Inc., 37 Kirkwood Road, Brighton, MA 02135.

Readers are encouraged to send to the Section Editor copies of reviews of statistical software appearing outside of the general statistical literature—for example, reviews appearing in journals such as *Journal of Applied Econometrics* or *Computer Applications in the Biosciences*.

Readers are invited to make two types of submissions to the "Editor's Notes" area of this Section: (1) specific software features that are currently unusual implementations; this must not be controversial and should have general appeal; examples appear in the November 1991 and February 1992 issues; please send a brief explanation of the statistic, a brief statement of its importance, and a small number of citations; and (2) views regarding general attributes or design issues for statistical software; an example appears in the November 1991 issue; other issues include: documentation, user interface, test data sets, forms of data editing allowed, output format control, speed, and so forth. Both types of submissions should be made directly to the Section Editor, Richard Goldstein, and are subject to extensive rewriting.

Findings and opinions expressed in every review are solely those of the author. They should not be construed as reflecting endorsement of the product, or opinions held, by *The American Statistical Association*, nor is any warranty implied about any product reviewed.

Editor's Notes

RICHARD GOLDSTEIN, Section Editor

This issue includes three reviews: two of the reviews cover "computer algebra systems"; Baglivo writes about Maple and Mathematica, and Nash writes about Derive; in the third review, Lurie covers some of the first Windows-specific packages.

Computer algebra systems are used to obtain symbolic solutions to mathematical problems. Stated this way, the use of such software in statistics may not be completely obvious. Baglivo has carefully chosen several examples to show uses as a teaching tool, as a data analytic tool, and in doing research in mathematical statistics. Both Maple and Mathematica are large, powerful packages with extensive, integrated, symbolic, numeric, and graphical capabilities. They also have extensive simulation abilities and have, or are rapidly gaining, broad document management abilities. As such, each requires substantial resources. Derive has many of the same capabilities but is much more compact and less broad, especially with respect to the performance of complete projects. Baglivo used Macintosh versions of these packages, but they run on many other platforms as well. Nash reviewed the DOS version of Derive. There is one other Macintosh computer algebra system: Theorist; however, the vendors never responded to my invitation to participate, so it is not included.

A number of other reviews of these systems are cited by Baglivo, and readers might want to examine some of these for different perspectives. Two appeared too recently for her to include: Crow (1994a, 1994b).

As new "operating environments" (such as Macintosh or Windows) become the rage, many users, including many reviewers, become adherents, even advocates for these platforms. For every such environment, there are both advantages and disadvantages for statisticians, and there are both good and bad statistical software packages. For those of us who make a living analyzing data by using

 © *1995 American Statistical Association*