

## A graphical approach to the analysis of censored data

Robert Gentleman<sup>1</sup> and John J. Crowley<sup>2</sup>

<sup>1</sup> *University of Waterloo, Waterloo, Ontario, Canada, and* <sup>2</sup> *Fred Hutchinson Cancer Research Center, Seattle, Washington, USA*

**Key words:** box plots, censored data, curve-difference chart, graphical methods, Kaplan-Meier estimator, quantile-quantile plot, rank plot, regression, smoothing, survivor function

### Summary

The presence of censoring has hampered the graphical exploration of survival data. We present several graphical approaches to the analysis of such data here, many based on functionals of the distribution function and estimated using the Kaplan-Meier estimate of the distribution function. Topics covered include comparing two samples, comparing many samples, and regression.

### Introduction

Censored data arise in a variety of settings. For the most part the analysis of censored data has relied upon techniques which are not graphical in nature. This paper discusses graphical methods for the examination of censored data. First, the comparison of two samples will be considered. Such situations arise most frequently in the health sciences where two treatments are to be compared. The discussion of methods for covariates with several levels and continuous covariates will follow.

Much of this work is motivated by the papers of Cleveland and McGill [1,2]. There is some overlap with Gentleman and Crowley [3,4]; however, the intent is much different. As in Gentleman and Crowley, the proposed methodology often arises from the use of a censored data estimate of the distribution function. The plots are thought of as functionals of the distribution function and *estimated* using an estimate of the

distribution function. The search for graphical methods will be tempered by the maxim that good graphical methods are ones in which the message of interest is easily and accurately perceived.

If one simply wants to determine whether the survival times depend upon the covariate, there are a variety of tests available. Kalbfleisch and Prentice [5] and Cox and Oakes [6] give details of a number of parametric and nonparametric tests for a variety of hypotheses. However, if one is unsure of the distribution of the observed data, it would seem prudent to examine the data more thoroughly. If the data are not subject to censoring such an examination would be mainly graphical in nature, since such an analysis usually allows detection of unexpected relationships in the data. We propose that when the data are subject to censoring a graphical examination is also warranted; however, there are few techniques that have been proposed or used extensively. This deficiency seems to stem from the fact that it is often difficult to incorporate the censoring infor-

mation in a sensible fashion. One may not rely upon simply encoding the censoring information into the plot symbol, since for individuals that were censored we know only that the true failure time is larger than the observed censoring time, while for observed failures the exact time is known.

Assume that the observed data are:

$$[(t_i, \delta_i, z_i)]_{i=1}^n$$

where  $t_i$  is the minimum of the true failure time and the censoring time  $\delta_i$  is an indicator of whether failure was observed, and  $z_i$  is the possibly vector-valued covariate information. We shall assume that the censoring times and the survival times are independent given the value of  $z$ . Let  $F$  denote the distribution function of the survival times, so that  $F(x) = P(X \leq x)$ , and let  $S$  denote the survivor function for the survival times, so that  $S(x) = 1 - F(x)$ . For data that are subject to right censoring, the Kaplan-Meier [7] estimate of the empirical survivor function is given by:

$$\hat{S}(t) = \prod_{i: t_i \leq t} \left[ 1 - \frac{\delta_i}{r_i} \right]^{\delta_i}$$

where  $r_i$  is the number of individuals at risk at time  $t_i$ . The Kaplan-Meier estimator reduces to the empirical survivor function in the absence of censoring. We shall also be quite interested in estimating various quantiles such as the median. The  $q^{th}$  quantile of the distribution function  $F$  is defined to be the smallest number  $\xi$  satisfying  $F(\xi) \geq q$ . In terms of the survivor function the  $q^{th}$  quantile is the smallest number  $\xi$  satisfying  $S(\xi) \leq 1 - q$ .

Regression models for censored data are usually defined in terms of the hazard, or instantaneous rate of failure. Two regression models that are commonly entertained for right censored data are the proportional hazards model and the accelerated failure time model. For the proportional hazards model the hazard is modeled as:

$$\lambda(t; z) = \lambda_0(t_0) \psi(z)$$

where  $\lambda_0$  is an unspecified baseline hazard function and  $\psi(\cdot)$  is the regression function. For the accelerated failure time model the hazard can be written:

$$\lambda(t; z) = \lambda_0(t\psi[z])\psi(z)$$

where  $\lambda$  and  $\psi$  have similar interpretations as in the proportional hazards model. In both cases the most commonly used regression function is  $\psi(z) = \exp(\beta z)$ .

In the sequel, three data sets will be examined in detail. The first is a study of survival following insult with the carcinogen DMBA to mortality from vaginal cancer in rats [8]. There are two groups which differed by a pretreatment regimen. The second is from the Southwest Oncology Group (SWOG) study 8292 [9]. This trial was set up to study whether surgery in addition to radiation therapy was beneficial to individuals with a single metastasis to the brain. The study was not randomized. The third data set comes from Appendix 1 of Kalbfleisch and Prentice [5] and concerns lung cancer data collected by the Veterans Administration in the United States. The reader is cautioned not to draw conclusions about these data from the examples given herein. They are intended to be pedagogical rather than to constitute a thorough examination of the data.

## Comparing two samples

Often the covariate information simply records which of two treatment regimens the individual was assigned to. In this case the comparison is similar to the standard two sample problem with uncensored data. The question that is generally most important is whether the survival times in one group tend to be larger than the times in the other. Often one can argue that it is the ranks of the survival times that are important rather than the actual survival times. An easy comparison of




Figure 1. A) A rank plot comparing the two treatment groups for the rat data. Observed survival times are indicated by filled octagons, censored observations are indicated by asterisks, and the group medians are indicated by triangles. B) Survival curves comparing the two treatment groups for the rat data. Group 1 is indicated by the solid line and Group 2 by the dashed line. C) A curve-difference chart for the survival curves given in 1B. This indicates the vertical distance between the two curves.

ranks comes from the *rank plot*. All the survival and censoring times are combined and ordered from smallest to largest, the smallest receiving rank one and the largest receiving rank  $n$ . The data are first split by group and then by whether they were censored. The four groups are then plotted against the ranks as has been done in Figure 1A. The failure times are plotted close to each other while the censored individuals have been plotted to the outside. A line is drawn down the center of the plot to distinguish the groups. Finally the medians of the two groups are indicated using triangles. The data used are the rat

data, and the plot indicates that those in Group 2 have larger ranks and hence larger survival times.

Another way of comparing the two groups would be to estimate the survivor functions separately and then compare them. This is the technique currently in standard use. From such a plot it is easy to see whether one survivor function dominates the other, i.e. whether  $S_1(t) > S_2(t)$  for all  $t$  or possibly for all  $t \geq t^*$ . Figure 1B is such a plot for the rat data. While it is the vertical distance between the two curves that is of interest our attention is invariably drawn to the area between 200 and 250 days where the two curves

appear to be close. The closeness at this point is mainly in the horizontal direction and not in the vertical direction. Cleveland and McGill [1,2] indicate that the use of parallel curves which they refer to as *curve-difference* charts can be very misleading. If one is really interested in the vertical distance between the two curves then the vertical distance should be computed and plotted. Figure 1C is the curve-difference chart for the two survival curves in Figure 1B. In this plot it can be seen that the vertical distance between the two curves is quite substantial in the region in which they appeared to be close, around 235 days.

Examining a plot of the estimated survival curves based on the brain data, Figure 2A, it is obvious that one curve dominates the other. Those receiving both radiation therapy and surgery, Treatment 1, are represented by the solid line, while those receiving only radiation therapy, Treatment 2, are represented by the dashed line. In this case a curve difference-chart would be either strictly positive or strictly negative (depending on how the difference was taken). While it is easy to see that the estimated survivor function for Treatment 1 is always above the estimated survivor function for Treatment 2, other questions are not so easily answered. For example, how do the median or other quantiles compare? If these questions are of interest then one might consider using a box plot. Side by side box plots for the brain data are presented in Figure 2B. The construction of the box plots is detailed in the Appendix. From this plot it can easily be seen that the median survival time for the surgery group exceeds the upper quartile of the survival times for those receiving only radiation therapy. In addition group size can be encoded in the width of the box, as was done here (the width of the plot is proportional to the square root of the sample size, in this case 51 and 23).

Another graphical method that is often used is the quantile-quantile plot (QQ-plot). QQ-plots are

graphical functionals of the two survivor functions and hence can be constructed for censored data using the Kaplan-Meier estimators. QQ-plots are often used to compare observed quantiles to quantiles from a reference distribution or to compare the quantiles from two observed samples. Figure 2C is a QQ-plot comparing the quantiles for the two treatment groups from the brain data. A forty-five degree line has been added to make it easier to compare the quantiles. The fact that not all quantiles are estimated with equal precision can make it difficult to interpret such plots in many practical situations. Michael [10] has suggested a stabilized probability plot to overcome this difficulty. His suggestion can be applied to right censored data.

If the data are thought to come from an accelerated failure time model then the QQ-plot comparing quantiles from the two groups should be a straight line through the origin with slope equal to the *acceleration*. In this case applying a logarithmic transformation to the survival times acts to make the variance of the estimated quantiles more stable and should be employed before plotting. If the data are thought to come from a proportional hazards model then all of the quantiles should lie below the forty-five degree line as they do in Figure 2C. However, the requirements for a proportional hazards model are somewhat more stringent and hence this feature does not necessarily imply that a proportional hazards model is appropriate.

If the two groups are thought to follow a proportional hazards model then it is common to simultaneously plot  $\log(-\log[S_i(t)])$  against  $t$ . If the proportional hazards assumption is tenable these curves should differ by a constant amount. Again it can be argued that one really should examine a curve-difference chart. Figure 2D is a plot of  $\log(-\log[S_i(t)])$  for the brain data where, once again, Treatment 1 is represented by the solid line while Treatment 2 is represented by the dashed line. It can be seen that the assumption of a constant difference is not tenable. Your eye

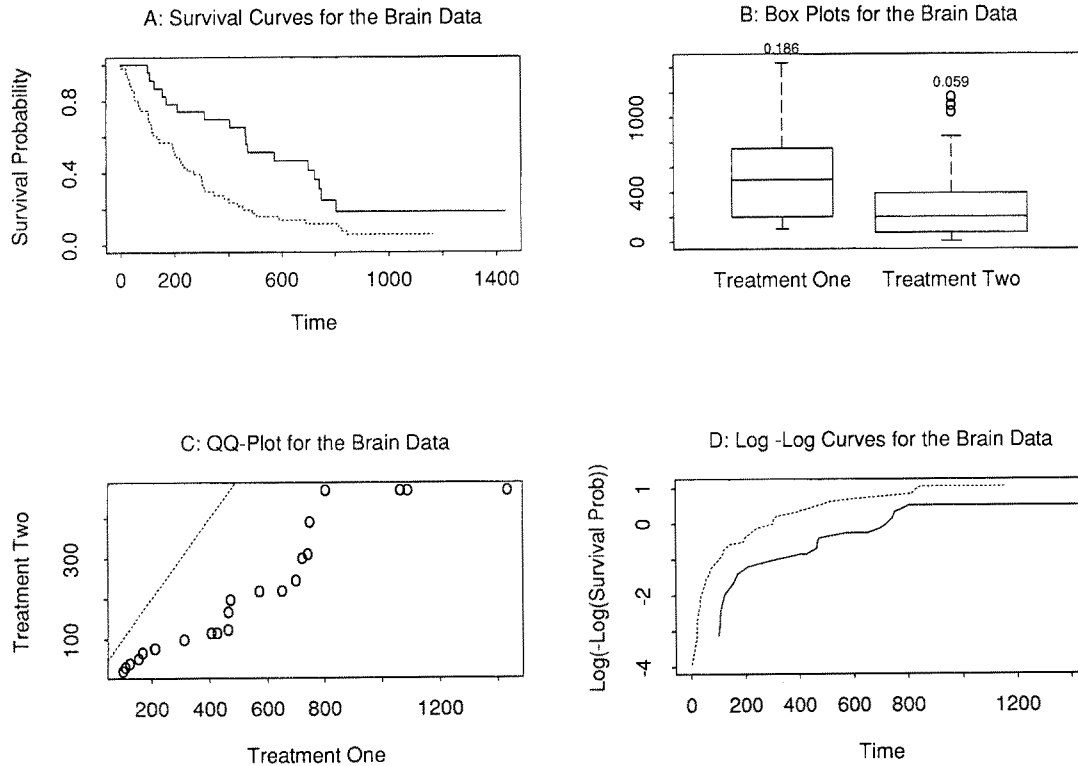


Figure 2. A) Survival curves comparing the two treatment groups for the brain data. Group 1 is indicated by the solid line and Group 2 by the dashed line. B) Box plots comparing the two treatment groups for the brain data. Sample size has been encoded in the width of the box, and the height of the survival curve for the last observation is given for both groups. C) A quantile-quantile plot comparing the two treatment groups for the brain data. The dashed line indicates where the points would lie if there were no differences in survival between the two groups. D) A plot of  $\log(-\log[S(t)])$  for the survival curves given in 2A. A constant vertical separation would indicate that the proportional hazards assumption was tenable.

wants to follow the two curves *around the corner* in the left portion of the plot. With some concentration one can see that the vertical separation decreases as survival time increases.

### Comparing many samples

In this section we discuss the use of graphical methods when the covariate is discrete but has more than two levels. Such covariates are often referred to as factors. We shall consider two types of factors, those for which there is an

ordering such as dose-level, and those for which there is no ordering such as the cell-type involved in lung cancer. These are referred to as ordinal factors and nominal factors, respectively. For nominal factors the methods are essentially those discussed for two sample comparisons, while for ordinal data one may also choose to employ a model which utilizes the additional information.

### Nominal data

For nominal data the level of the factor simply describes which group the individual is in, while

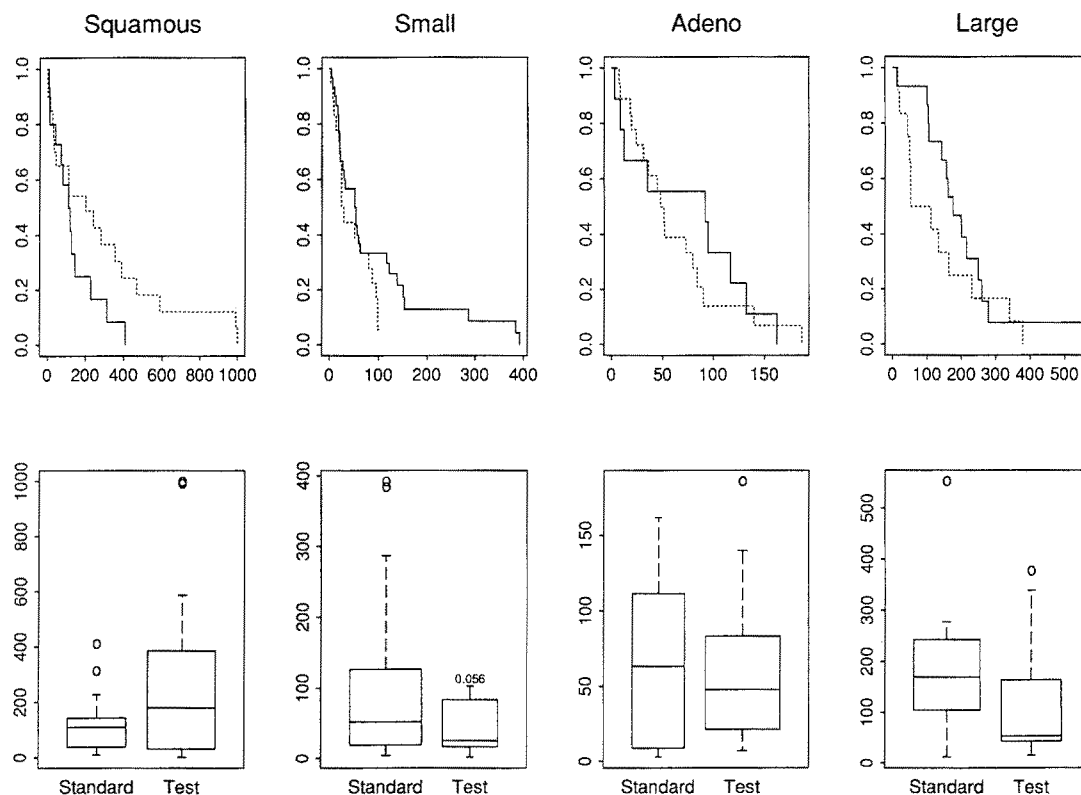


Figure 3. The top four frames provide survivor curves comparing the standard and the test treatments for each of the four cell types for the lung cancer data. The bottom four frames provide the same comparison via box plots. The time scale is not the same for all plots.

for ordinal data the level of the factor provides quantitative information. With the lung cancer data one factor is the cell-type involved in the cancer. This factor comes at 4 levels, but we have no preconceived notion about the relationship of the response for those with squamous cell involvement versus those with large cell involvement. Later we will divide the observations into groups based on their age yielding an ordinal factor.

A first approach is to compute the Kaplan-Meier estimate of the survivor function for each level of the factor separately and to then plot these estimates. Figure 3 compares the estimated survivor functions, in the top frames, with parallel box plots, in the bottom frames. While in general

one should use the same scales, both vertical and horizontal, we have not done so for all plots in this figure. Using the same time scale for all plots compacts the estimates for the Adeno group substantially and makes them difficult to examine, so we have chosen to use different scales. Comparisons down columns are relatively straightforward but comparisons across the rows cannot easily be made.

The problems faced in the comparison of two samples are magnified when comparing many samples. The plot of the estimated survival curves often looks completely incomprehensible. For the lung cancer data we are also concerned with comparing the standard treatment with the test treatment. This has been done in Figure 4,

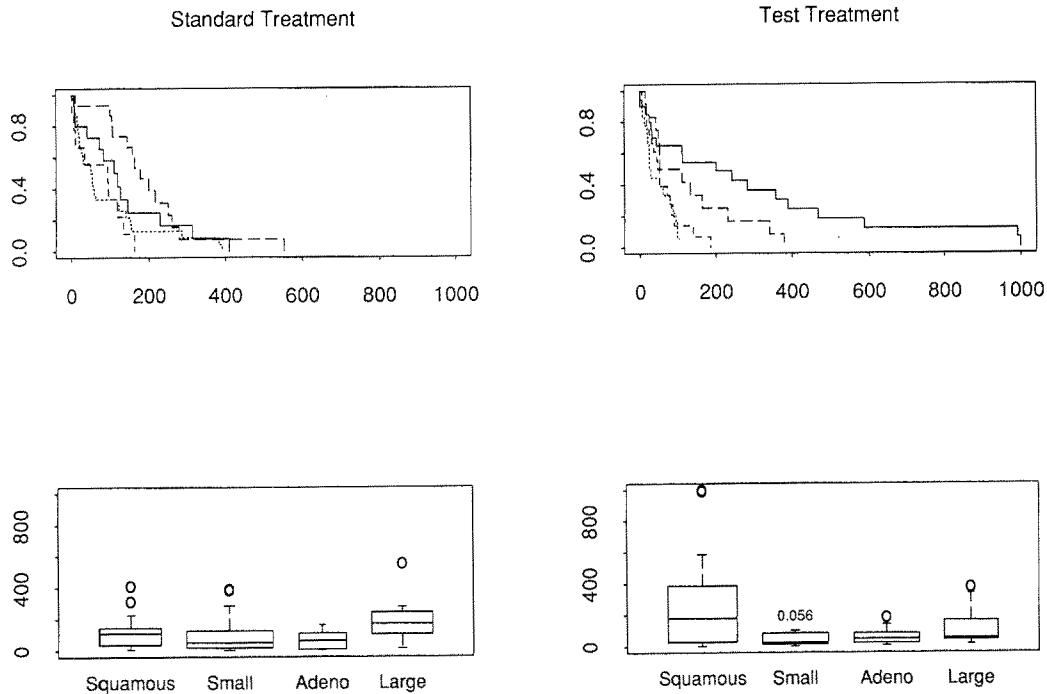


Figure 4. The top two frames compare the two treatment groups and each of the four cell types for the lung cancer data via survival curves. The bottom two frames use box plots to provide the same comparisons. Sample size has been encoded in the width of the box.

the top two frames of which contain estimates of the survivor function while the bottom two contain box plots for the same data. One cannot easily make the visual comparison of the 4 survival curves. Great care and concentration are necessary to extract the relevant information, while for box plots the information that is provided is more easily extracted. These plots indicate that the test treatment appears to be quite advantageous for individuals with squamous cell involvement but has little, if any, effect on survival for the other cell types.

The problem seems to be that even when there are as few as 4 survival curves being plotted it is very difficult to accurately assess their inter-relationships. Plotting all pairwise comparisons separately may be useful but it is difficult to get an overall impression. Any of the plots discussed

in the previous section, such as QQ-plots, could be used to make pairwise comparisons.

### Ordinal data

When the factor is ordinal there are several options that can be exercised. The factor may be analyzed with little regard to the ordering using the techniques discussed above, it can be used as a covariate in a regression model, or it can be used to define strata. Both the proportional hazards model and the accelerated failure time model can be extended by the incorporation of strata. While some covariates should be included in the model as regressors, other variables do not have the appropriate relationship with the response for them to be included in this manner.

For example, it may be that the baseline hazard is different at different levels of the factor and hence a proportional hazards model cannot be used. However, if there are other covariates involved then this factor, let us say  $z$ , could be used to define strata. For such models the standard assumption is extended to the following:

$$\lambda(t; x, z) = \lambda_{0,z}(t) \psi(x)$$

with a similar modification of the accelerated failure time model.

As mentioned previously, a plot of  $\log(-\log[S(t)])$ , where  $i$  indicates the level of the factor, can be used to determine whether the hazards are proportional, since on this scale proportional hazards result in curves which exhibit approximately constant separation. Again, the argument that this determination should be made from a curve-difference chart can be made. Further, the relatively lower precision of the estimated survival curve for large values of  $t$  may be problematic.

Examining the data at each level of the factor separately can often be informative. For example plotting the survival curves may indicate a trend towards higher or lower survival times for increasing levels of the factor. Using the ordered levels of the variable to plot various graphics can reveal whether or not there is an effect on the survival distribution which depends on the level of the ordinal variable in a monotonic manner.

### Continuous covariates

When the covariate is continuous it is somewhat more difficult to ascertain the nature of the relationship with survival time. There are two major difficulties that are encountered. Standard regression models generally state the distribution of response given the value of the covariate. In linear regression this is usually taken to be the Normal distribution and hence is completely described by its mean and variance. The situation

is not quite so simple for censored data regression models. For example, the proportional hazards model has an infinite dimensional nuisance parameter (the baseline hazard function). The second part of the problem is that the presence of censoring, particularly if it is heavy, can disguise the true form of the relationship between the covariate and survival time.

However, there are a few techniques which have proven quite useful when exploring censored data. One of these is discretization [11], that is, the continuous covariate is transformed into a discrete covariate by binning all those observations with similar values of the covariate into a group. The data can then be analyzed using the methods for ordinal data of the previous section. This method has the advantage that if the bins are formed so that each observation appears in only one bin then the results are independent, in the statistical sense. A second method is to employ *smoothing* techniques. Scatter plot smoothers were developed in order to enhance the detection of relationships in scatter plots. Such detection is often adversely affected by the presence of outliers or by large numbers of observations. On the simplest level one can think of a scatter plot smoother as an estimate of the mean of the  $y$ -variable conditional on the corresponding value of the  $z$ -variable. Since with censored data one often does not observe the mean, a different measure of the center will be needed. A candidate for this is the conditional median. In fact, plotting various running quantiles can vastly improve our visual perception of the relationships in the data.

For the first suggestion one needs to decide how to bin the data. The manner in which this is done can often be important; there are many examples of histograms that change shape dramatically when slightly different binning algorithms are used. However, we suggest using the rather simplistic approach of dividing the data into approximately five groups of roughly equal size. Some caution must be exercised if there are re-






Figure 5. Box plots comparing the survival time distribution for different age groups. The data were binned into non-overlapping groups.

peated values in the covariate, as all individuals with the same covariate value should be in the same group. If this method is employed in conjunction with the smoothing methods we discuss next, any serious anomalies should be readily detected. Since the data have been divided up on the basis of the covariate they may be examined using the methods for ordinal data discussed in the previous section. Figure 5 provides parallel box plots using the age covariate from the lung cancer data set of Kalbfleisch and Prentice [5]. The data have been split into 5 groups via the following rules:  $age < 48$ ,  $48 \leq age < 60$ ,  $60 \leq age < 64$ ,  $64 \leq age < 68$ , and  $age \geq 68$ . These splits were chosen so that there were approximately 28 observations in each group. As we can see in Figure 5, there appears to be little relationship between age and survival time, with perhaps a slight increase in survival for those individuals in the 48 to 59 year age range.

### Smoothing

In this section we shall discuss two types of smoothing, one nonparametric and the other model based. Scatter plot smoothers have long been recognized as advantageous when there are anomalies that may affect our ability to accurately perceive the relationship between the two vari-

ables being plotted. As the presence of censoring does this, one can expect that the use of smoothers should be quite helpful. As discussed above we shall consider smoothers which provide estimates of the quantiles of the distribution of survival time conditional on the value of the covariate. Gentleman and Crowley [3,4,12] discuss the issues involved in somewhat more detail.

A basic premise of smoothing is that the conditional distribution of the  $y$ -variable changes slowly with respect to changes in the  $x$ -variable. Thus when estimating this distribution for a particular value of the  $x$ -variable it is reasonable to increase the sample by using not only those observations with that value of the  $x$ -variable but also those observations with similar values. Sometimes schemes which weight the contribution of each observation depending on the distance from the target value are used. The number included in each computation is often referred to as the span; this is often considered as a proportion of the total number of observations. We have chosen to use symmetric  $k$  nearest neighborhoods for our calculations. In this case span selection devolves to choosing an appropriate value of  $k$ . At each distinct value of the covariate the  $k/2$  nearest neighbors to the left of the target value and the  $k/2$  nearest neighbors to the right form the neighborhood. When the target value is so close to either the right or the left boundary that there are not  $k/2$  neighbors in both directions one simply uses the ones, if any, that are there. Most smoothers do not provide reliable estimates near the boundary of the covariate, and caution should be used in interpreting effects seen near its edges.

Smoothing exchanges the issue of choosing bins, discussed above, for the problem of choosing the bandwidth, or smoothing parameter. There are several reasons why one would like to have an automatic choice of smoothing parameter. Unfortunately this issue has not been satisfactorily addressed for censored data, and no automatic method exists. It is important to ensure that there are sufficient observed failures in every interval in

order to adequately estimate the survivor function or quantiles, as required.

For nonparametric smoothing one proceeds as follows. For each distinct value of the covariate those observations in the symmetric  $k$  nearest neighborhood are determined. Using only those observations a Kaplan-Meier survivor function is computed, and using this the desired quantiles are computed. These are then stored and used as the estimates of those quantiles for that value of the covariate. The estimated quantiles can then be used to add a conditional quantile line to a plot of the data. This method was used on the covariate age from the lung cancer data, and the result is plotted in Figure 6. Three conditional quantiles, the 0.25, 0.5, and 0.75, were computed. There are two sets of conditional quantile lines plotted in this figure, and those with the shorter dashes correspond to the procedure described above. We have provided two sets based on different spans. For a span of 0.3 approximately 41 observations were used in each neighborhood, while for a span of 0.6 the number was approximately 82. Notice that for the span of 0.3 there is some indication that the upper quartile is elevated in the region from 45 years to 60 years, in agreement with the impression given by the discretization method. For a span of 0.6 the estimated conditional quantiles are pretty much constant, indicating little dependence on age.

The model based smoothing method used employs the proportional hazards model and is discussed more fully in Gentleman and Crowley [4]. In this case again one uses local information, but rather than estimating a survivor function in each neighborhood a proportional hazards model is fitted to the data. For this model, the baseline hazard is assumed to be common to all individuals and is estimated using all the data, while the covariate effect is estimated locally. This model was also fitted to the lung cancer data, using the same span size, and the results used to estimate the same three quantiles as above. The corresponding conditional quantile lines were

plotted, using long dashes, in Figure 6. Having both sets of conditional quantiles on the same plot facilitates the comparison of the results from the two different methods. Using a span of 0.3 there seems to be some substantial discrepancy between the two methods, at least with regard to the level of the survivor function, but the overall shape is approximately the same. With a span of 0.6 the agreement between the two methods is quite good.

The conditional quantiles can also be used to determine the appropriateness of certain models. For both the proportional hazards model and the accelerated failure time model, the usual form of the covariate effect is  $\psi(z) = \exp(\beta z)$ . In this case a straightforward calculation (see Gentleman

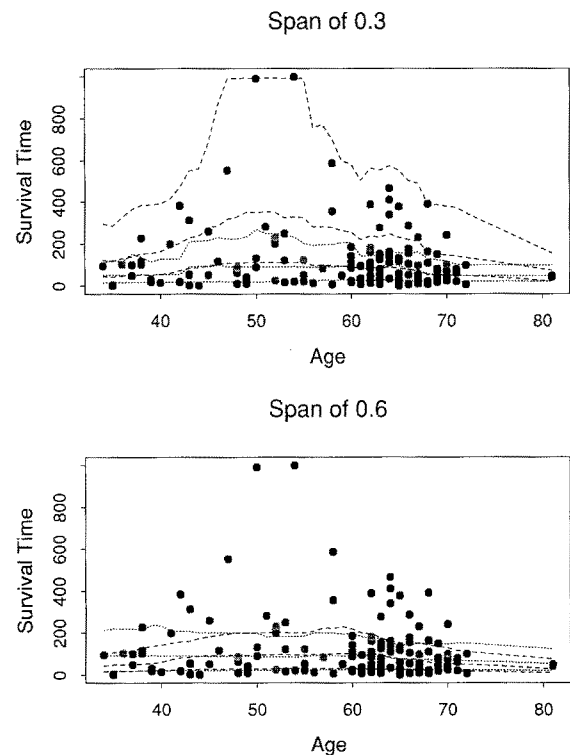


Figure 6. Smooth estimates of the quantiles (0.25, 0.50, and 0.75) for survival time as a function of age are plotted. The long-dashed lines are estimates based on a proportional hazards assumption, while the short-dashed lines are based on a non-parametric method.

and Crowley [4]) indicates that in general,  $\beta > 0$  implies that the conditional quantiles when considered as a function of  $z$  need to be monotonically nondecreasing for the accelerated failure time model and monotonically nonincreasing for the proportional hazards model. The converse is true if  $\beta < 0$ . As the conditional quantiles in Figure 6 do not appear to behave in a monotonic fashion, one would tend to believe that a covariate effect of the form  $\psi(z) = \exp(\beta z)$  would not be appropriate in this situation.

## Discussion

We have demonstrated a number of graphical methods that can be used for the analysis of data that are subject to censoring. These suggestions provide a variety of means of examining the data, both with regard to data integrity and with regard to model checking. Thus, the user can fairly easily get a reliable impression of the relationships that exist between various covariates and survival time prior to fitting any model. Such an examination, exploratory data analysis, is often a very valuable tool which can lead to a more appropriate choice of model or new insight into the data. In addition to providing several explicit descriptions of graphical methods for censored data, we have provided an algorithm, of sorts, for constructing new ones. New graphical methods can be developed by using the Kaplan-Meier estimate, or any other estimate, of the survivor function in plots originally developed for data not subject to censoring.

## Acknowledgements

The data from SWOG 8292 are used with the permission of Dr. Harmon Eyre, Chairman of the Southwest Oncology Group Brain Tumor Committee, and Dr. William Sause, who was in charge of the study. We would also like to thank Dr.

Charles Coltman, Chairman of the Southwest Oncology Group, for the encouragement to use the Group's data, and J.A. McDonald for many helpful discussions. The first author's research was supported by a Natural Science and Engineering Research Council of Canada grant.

## Appendix

### *Details of the Box Plot Construction*

For a given sample compute the Kaplan-Meier estimate and use it to find the median as well as the upper and lower quartiles. If the censoring is not too heavy these will all be observed and one proceeds in essentially the same manner as with uncensored data. The interquartile range (IQR) is the upper quartile minus the lower quartile. The quartiles are used to form the box and the median is drawn in as a line across the box at the appropriate position. The lower whisker extends below the box to the smallest observed failure time within a distance of 1.5 times the IQR of the lower quartile. The upper quartile extends from the upper quartile to the largest observed failure time within a distance of 1.5 times the IQR. Observed failure times further away from the center of the data than the limits on the whiskers are plotted individually. The value of the survivor function at the largest observation is plotted at that point, if it is not zero.

If not all of the quartiles are observed, then plot the observed ones and extend the sides of the box to the largest observed failure time. At this point plot the value of the survivor function centered between the sides of the box. Since the upper quartile has not been observed the IQR will need to be estimated. This can be done via linear interpolation, but some attention should be paid to the effect that a decidedly asymmetric distribution can have. The number of observations in each group can be encoded in the width of the box.

Computer routines to perform this and most of

the other methods described herein, implemented in the S language, are available from the first author.

## References

1. Cleveland WS, McGill R: Graphical perception: theory, experimentation, and application to the development of graphical methods. *J Am Stat Assoc* 79:531-554, 1984.
2. Cleveland WS, McGill R: Graphical perception: the visual decoding of quantitative information on graphical displays of data (with discussion). *J Royal Stat Soc [A]* 150:192-229, 1984.
3. Gentleman R, Crowley J: Graphical methods for censored data. *J Am Stat Assoc* 86:678-683, 1991.
4. Gentleman R, Crowley J: Local full likelihood estimation for the proportional hazards model. *Biometrics* 47:1283-1296, 1991.
5. Kalbfleisch JD, Prentice RL: *The Statistical Analysis of Failure Time Data*. John Wiley and Sons, New York, 1980.
6. Cox DR, Oakes D: *Analysis of Survival Data*. Chapman and Hall, London, 1984.
7. Kaplan EL, Meier P: Nonparametric estimation from incomplete observations. *J Am Stat Assoc* 53:457-481, 1958.
8. Pike MC: A method of analysis of a certain class of experiments in carcinogenesis. *Biometrics* 22:142-161, 1966.
9. Sause WT, Crowley JJ, Rotman M, Mowry PA, Bouzoglou A, Borst JR, Selin H: Solitary brain metastasis: results of an RTOG/SWOG protocol evaluation. Surgery and radiation therapy versus radiation therapy alone. *Am J Clin Oncol* 13:427-432, 1990.
10. Michael JR: The stabilized probability plot. *Biometrika* 70:11-17, 1983.
11. Cox DR: Some remarks on the role in statistics of graphical methods. *Applied Statistics* 27:4-9, 1978.
12. Gentleman R, Crowley J: Smoothing censored data. Technical Report 90-13, University of Waterloo, Dept. of Statistics and Actuarial Science, 1990.