

A gentle introduction to survival analysis

Angelo Tinazzi, SENDO-Tech S.r.l., Milan, Italy
Martin Scott, Numerus Ltd, Tuebingen, Germany
Anna Compagnoni, SENDO-Tech S.r.l., Milan, Italy

ABSTRACT

Time-to event analyses are commonly used to analyze a clinical trial. Survival Analysis is a necessary tool to analyze the time to the occurrence of an event of interest, which can be either adverse, positive or neutral. Analysis of survival data requires special techniques because the event of interest does not necessarily occur for all patients before the end of the study. Survival methodology has the advantage of being able to use all available data to estimate the survival probability, including data derived from patients who did not experience an event.

This paper is a tutorial addressed to statistical programmers who are required to produce derived datasets containing time-to-event endpoint variables and summary tables containing survival data estimates (e.g. median survival). The paper is composed of two main sections.

In the first part, the main elements of survival analysis theory will be introduced:

- the Kaplan-Meier method for survival probability estimation;
- the log-rank test;
- the hazard ratio calculation;
- the interpolation of a Cox regression model.

In the second part, the main programming aspects of survival analysis will be discussed:

- creating derived datasets containing time-to-event endpoints;
- working with longitudinal data (repeated measures, such as follow-up);
- the use of the main SAS[®] procedure for survival analysis (e.g. proc LIFETEST and proc PHREG);
- extracting the most relevant information (e.g. median survival) from a SAS output using the ODS for creating ad-hoc summary tables.

INTRODUCTION

In many clinical studies, either experimental or observational, the time until the occurrence of an event is the main outcome of interest; such an event can be either adverse, such as death or disease progression or recurrence of a tumor, or positive, such as birth or disease remission, or neutral, such as cessation of breast feeding or the completion of a task. In all these cases, it is conventional to talk about survival endpoints; the term survival may be misleading because the techniques are applicable to any well-defined event although traditionally death was the event of interest. Events in survival analysis are defined by a 'transition' from one discrete state to another at an instantaneous moment in time.

The typical feature of this kind of data is that not all the subjects will have experienced the event at the end of the follow-up period. This is called censoring, meaning that the observation period ended without observing the event of interest, either because the subject has not yet experienced the event by the end of follow-up period, or because the patient is lost to follow-up, or because the subject experiences a different event that makes further follow-up impossible.

It is not possible in these cases to know if and when the subject would have experienced the event and we refer to this as right censoring meaning that a subject's follow-up stops before observing the event; while left censoring refers to when the event has occurred before the subject has entered the study.

Classical statistical methods, such as T-test or Chi-square Test, are not appropriate to analyze survival data because they are not conceived to take censored data into account. All subjects (including those with censored observations) provide useful information for the trial, and should therefore not be excluded from the analysis. On the other hand, considering the censoring time as a true survival time for these subjects would lead to a biased estimate of their survival time and probability. For these reasons there is the need to use a set of specific statistical methods to handle survival data.

To make an example, consider the following two clinical study endpoints:

- the risk of dying in Intensive Care Unit (ICU) within 30 days from the admission;

PhUSE 2008

- survival rate of advanced breast cancer patients at 5 years from the diagnosis date.

In the first case you can describe the endpoint results by using a simple % rate (events/totals), and eventually using the chi-square test, for two main reasons:

- the limited length of required observation period (30 days): there is no risk for losing patient;
- at the end of the 30 days period, a patient can be either dead or alive (discharged or not).

In the second case it is more appropriate to use a survival analysis approach for the following reasons:

- with a required length of follow-up (5 years), there is an high risk to loose patient (lost to follow-up);
- usually this kind of study requires a longer accrual period (sometimes varying from 2 to 5 years), therefore if you perform an interim analysis not all the patients had enough time to reach (alive) the 5 years survival endpoint;
- the time when the event has been observed (e.g. 1 year vs 3 years) with this “chronic” disease, make a substantial difference, therefore simply checking how many patients died in the 0-5 years period and eventually removing lost to follow-up patients, it is not appropriate.

INCOMPLETE DATA

In clinical practice not all the subjects are enrolled at the same time and thus the follow-up period can vary from one subject to another. For these reasons it is suitable to convert calendar-time data to time-to-event data, which is a more appropriate format to perform survival analysis.

The starting point is set to the same (e.g. entry into study or diagnosis time, will be the beginning point of the study defined as time=0) for all enrolled subjects, and all the possible outcomes are coded depending on the event of interest.




Figure 1: Example of Survival Data.

R = Relapse; D = Death from ovarian cancer; Do = Death from other cause;

A = Alive; L = Loss to Follow-up; x = Death; □ = Censored

PART I: STATISTICAL METHODS

MAIN ASSUMPTIONS

To ensure a correct estimate of the survival probability in presence of censored data, three main assumptions are necessary.

First, we assume that subjects who are censored have the same survival distribution as those who continue the observation during the study: that is to say, the censoring process is independent of the primary endpoint. If the reason for censoring is due to negative treatment effects or to the worsened subject condition, this censoring would lead to a biased estimate of the survival probability.

The second assumption derives from the fact that subjects who first joined the study have a longer observation period and thus have more chances to experience the event than subjects enrolled at the end of recruitment period. We assume that survival probability is the same for subjects enrolled at the beginning and at the end of recruitment period.

For some kind of events of interest, like tumor recurrence or progression, it is not possible to identify the exact date

PhUSE 2008

when the event happens, because the only information we would have is the event happened between two subsequent examinations. In this case we need to assume that the event happens at the time it is detected, otherwise we would obtain a biased estimate of the survival times.

CALCULATION OF SURVIVAL PROBABILITY: THE KAPLAN-MEIER METHOD

Survival data are described in terms of two probabilities: survival and hazard.

The first one is the probability for an individual to survive from the starting time until a specified future time t .

To estimate the proportion of subjects surviving at a given time point, and hence the survival probability to that time for the generic population from which the sample is extracted, the Kaplan-Meier method [1][2][3][4][5][6][7], also called product-limit estimator, is commonly used, which allows to deal with censored information.

The simple idea underlying this method is that to survive to k intervals from the start of the study, it is necessary to survive to each one of the previous intervals, and then also to the k^{th} one. This principle allows working with conditional and cumulative probabilities.

The conditional survival probability for the generic time interval (on the condition that the subject was survivor at the beginning of the interval and then to all the previous) is given by:

$$p_i = \frac{r_i - d_i}{r_i}$$

where r is the number alive and at risk at the beginning of the i^{th} interval, and d the number of failures during the same interval. Observations censored at any given time influence the number at risk at the start of the subsequent time interval.

Survival at any time point is calculated as product of the conditional probabilities of surviving each previous time interval.

$$S(t) = \prod_{i \leq t} \left(1 - \frac{d_i}{r_i}\right)$$

Figure 2 is an example of survival probability calculation, derived from a SAS output referred to time to progression data (time expressed in weeks).

Product-Limit Survival Estimates					
Time	Survival	Failure	Survival Standard Error	Number Failed	Number Left
0.0000	1.0000	0	0	0	17
6.0000*				0	16
10.0000				1	15
10.0000	0.8750	0.1250	0.0827	2	14
16.0000	0.8125	0.1875	0.0976	3	13
16.0000*				3	12
16.0000*				3	11
17.0000	0.7386	0.2614	0.1133	4	10
20.0000	0.6648	0.3352	0.1237	5	9
23.0000	0.5909	0.4091	0.1302	6	8
26.0000	0.5170	0.4830	0.1332	7	7
26.0000*				7	6
29.0000*				7	5
30.0000*				7	4
31.0000*				7	3
38.0000*				7	2
43.0000	0.2585	0.7415	0.1946	8	1
44.0000*				8	0

NOTE: The marked survival times are censored observations.

PhUSE 2008

Figure 2: Survival probabilities Calculation

At week 0, all enrolled subjects (17) are alive and at risk to experience the event of interest (progression). Survival probability is thus set to 1¹.

At week 6, a censoring occurs (but number failed is still 0): it does not modify the survival and failure probabilities, but the number of subjects at risk for the subsequent interval decreases by one (number left = 16).

At week 10, two events occur, thus all the statistics are modified: the number failed increases by two, while the number left still at risk for the subsequent interval reduces to 14. Survival probability for the interval [time 0 to time 10] is then equal to $(16-2)/16=0.8750$.

A new event occurs at week 16, together with two censored observations; the survival probability for the interval [week 10 to week 16] is equal to $(14-1)/14=0.9286$. To calculate the cumulative probability for a subject to survive from the start of the study until the week 16 it is now necessary to cumulate the survival probabilities for each time interval. Cumulative survival probability until time 16 is computed as $0.9286 \times 0.8750 = 0.8125$.

At week 17 a new event occurs: the survival probability for the time interval [week 16 to week 17] is computed as $(11-1)/11=0.9091$.

Cumulating all the previous single probabilities, we can easily obtain the probability for a subject enrolled at the beginning of the study to survive without experiencing the event until week 17: $0.9091 \times 0.9286 \times 0.8750 = 0.7387$.

The same process can be followed until week 44 when no more events are available and the survival probability can't be modified anymore.

SURVIVAL PROBABILITIES CALCULATION: THE LIFETABLE METHOD

In some situations it is not possible to know or to record the exact time when the event occurs; the only available information is that the event occurred in a certain time interval. In this case the survival probability cannot be updated every time an event happens, but once in a fixed time unit: the derived survival curve and probability are called life-table or actuarial estimates.

With this method, the whole period of observation is divided into a series of time intervals, where the basic assumption is that censorings are uniformly distributed in each interval, so that the average number of subjects at risk during each interval is given by:

$$n'_j = n_j - \frac{c_j}{2}$$

Where n'_j is the number of subjects alive and at risk at the beginning of the time interval j and c_j is the number of censored observations in the same interval.

As previously seen with the Kaplan-Meier method, the cumulative probability to survive until the time t is given by the product of the probabilities to survive to each previous time interval and also to the last one including t .

$$S(t) = \prod_{j \leq t} \left(1 - \frac{d_j}{n'_j}\right)$$

d_j is the number of deaths or events observed in the time interval j .

THE GRAPHICAL REPRESENTATION

Survival probability can be plotted against survival time, obtaining a survival curve: the survival curve is a non-increasing function with a step corresponding at each time when an event is observed; the step is used since the estimated survival curve remains at a plateau between successive patient event time and it drops instantaneously at each time of event to a new level. Times when a censoring happens are usually indicated with a small circle or a short vertical line. The plot will never reach '0' if the patient with longest observed time has in fact died. Were such a patient still alive then the curve would have a plateau commencing at the time of the last event and continuing until the censored observed time of this longest 'surviving' patient.

In certain circumstances a graph of $1-S(t)$, rather than $S(t)$, is plotted against time to give the cumulative 'event' curve.

From the graphical representation of a survival curve it is possible to extract relevant summary statistics, mainly the median survival time and the survival probability at a certain time point.

The median survival time is defined as the time point when half of the subjects have experienced the event of interest, and can be easily determined by drawing a line (the red line in Figure 3) starting from the value 0.5 on the vertical axis until you reach the survival curve, and then from the curve till the horizontal time axis. The exact value of the median

¹ It can be also expressed as % (see figure 6)

PhUSE 2008

survival time can be calculated with a proportion involving the length of the axis and the maximum follow-up time observed. Similarly you can extract the median from the estimates (figure 2) by searching for the first estimate below 0.50; in the case the 0.50 (or below) is not reached, we say “the median of survival has been not yet reached”. This latter situation, can happen when cohorts with low event rate or in the case of early (interim) analysis. Again in the same way you can extract other percentiles of interest, typically the 25% and the 75%, by searching the point in which the curve reach the point of interest (e.g. .25 and .75).

The survival probability at a certain time point (e.g. 20 months, Figure 3) can be calculated in a similar way: as shown by the blue line in Figure 3, it is necessary to draw a line from the chosen point of the horizontal time axis to the survival curve, and then until the vertical survival probability axis. Once again, the exact probability can be deduced using a proportion involving the length of the vertical axis and the maximum value of the survival probability.

It is difficult to judge precisely when the right-hand tail of the survival curve becomes unreliable. However, as rule of thumb, the curve can be particularly unreliable when the number of patients remaining at risk is less than 15 or more in general less than 5%.

Another important statistic you may be required to provide, is the median of follow-up. The median of follow-up is an indicator of how ‘mature’ are your data by giving an estimate of the length of the observation period (e.g. how many months on ‘average’ the patients have been followed since they have been enrolled). To calculate the median of follow-up we can use survival analysis approach by censoring patients who experienced the event. The reason for that is because you don’t know how long you could follow a patient if this did not experience the event.




Figure 3: Survival Curve with Censored Observations: Median Survival (red line) Time and Survival Probability at a Fixed Time Point (blue line).

COMPARING SURVIVAL CURVES

An interesting issue in the clinical practice is the comparison of the survival functions of two (or more) groups of subjects.

Classical statistical methods such as T-test or Chi-square Test are not appropriate to compare the proportion of survivors for the reasons described before.

Alternatively, one could compare the proportions surviving at one or more pre-specified time points, but this approach ignores the total survival experience of the groups during the whole follow-up period. Moreover, the time points are arbitrary chosen.

The most common method used to compare survival curves is then a statistical hypothesis test called log-rank test [8], which takes the whole follow-up period into account and does not require any assumption about the distribution of the survival function, which is rarely normally distributed and is often skewed because it often comprises many early events and only a few late ones.

The null hypothesis for the log-rank test is that there is no difference between the survivals of two or more populations that are being compared (i.e. the probability of the event of interest occurring at any time point is the same for each

PhUSE 2008

population).

The comparison is based on the difference between the observed number of events in each group and the expected number of events in case of non-difference between the two groups.

$$\chi^2_{\log-rank} = \sum_g \frac{(O_g - E_g)^2}{E_g}$$

where O is the number of observed events in each group g , and E is the total number of expected events in each group g . O and E are calculated for each time when an event happens; if a survival time is censored, then the subject is considered to be at risk during the interval of censoring, but not anymore for the subsequent intervals.

The test statistic is then compared with a χ^2 with $g-1$ degrees of freedom.

The log-rank test assigns equal weight to each event at whatever time it occurs.

If researchers believe that differences between certain parts of the survival curves being compared are of greater interest than others, then different weights can be assigned to the deaths, depending on the time when they happen. For example, if it were known that a new treatment is helpful to avoid early mortality, then it would be useful to compare the first part of the survival curves, without paying too much attention to the subsequent time intervals.

The best way to achieve this aim is to construct a statistical test using as weights the number of patients at risk at each time point: this test is a variant of the log-rank test and it is known as Gehan or generalized Wilcoxon or Breslow test and can be calculated as follows:

$$\chi^2_G = \sum_{g,t} \frac{R_t (O_{gt} - E_{gt})^2}{R_t^2 E_{gt}}$$

where R is the number of patients at risk in each group and for each time point.

Note that these are tests of significance, and do not provide an estimate of the size of the differences in survival probabilities between the groups.

COX REGRESSION MODEL AND HAZARD RATIO

While with previous described methods, either the Kaplan-Meier or the Life-Tables method, the survival data are analyzed using an univariate approach.

To approach survival analysis in a multivariate way, the Cox model is used [9][10]. The Cox model a semi-parametric regression model used to test the effect of a set of covariates $X_1, X_2, \dots, X_n$ on the time-to-event variable; the model for the hazard function is the following:

$$h(t) = \lambda(t) \cdot \exp\{\beta_1 x_1 + \beta_2 x_2 + \dots + \beta_n x_n\}$$

where $h(t)$ represents the hazard function and $\lambda(t)$ is an unspecified initial hazard function.

The hazard is a function of time and represents the instantaneous probability that the event of interest will occur at a specified time t given that it has not occurred prior to time t .

Hazard is defined as follows:

$$\lim_{\Delta t \rightarrow 0} \frac{\text{observed events}(t) / N(t)}{\Delta t}$$

Where N is the number of subjects at risk at the beginning of the time interval.

The Cox model (that, as with the Kaplan-Meier method, also adjusts for censored data) can include one or more numeric covariates (which can be also time-dependent) and does not make any assumption on the distribution of the event times.

The main assumption of the Cox proportional hazard regression model is that the hazard ratio between the treatment groups remains constant over time, even if, as it is reasonable to assume, the hazard rate does change over time.

D. R. Cox, after whom the model is named, developed a method to estimate the coefficients β_i based on a maximal partial likelihood, which is a modification of the maximum likelihood method.

Chi-square tests are performed to evaluate the null hypothesis of non-significance of the β_i parameters. A positive parameter estimate indicates that the hazard increases with increasing values of the covariate, while a negative estimate shows that the hazard decreases with increasing values of the covariate.

The magnitude of the effect of a specific covariate on the event time is often expressed as hazard ratio [11]: this is an estimate of the relative survival experience of two groups and can be calculated as

PhUSE 2008

$$\frac{O_e / E_e}{O_c / E_c}$$

If the hazard ratio indicates a beneficial effect of the treatment on study, this implies that the time to the observed endpoint was modified by treatment.

The hazard ratio derived from the Cox regression model does not translate directly into information about the duration of time until the event occurs, but it can be converted, anyway, in terms of absolute risk reduction, showing the difference between the risk in the control group (set to 1) and the risk in the treated group (that is hazard ratio itself). Absolute risk reduction is defined thus as $1 - \text{HR}$, and can be expressed either as a probability or as a percentage.

PART II: STATISTICAL PROGRAMMING ASPECTS

ILLUSTRATIVE STUDY: OVARIAN CANCER STUDY

This data set refers to a multicenter, phase III, randomized clinical trial involving women with epithelial ovarian carcinoma randomized to either receive a more (Lymphadenectomy) or less (No Lymphadenectomy or Surgery Only) invasive surgery; the main end-points of the study included progression-free and overall survival [12].

427 eligible patients were enrolled: 211 were randomly assigned to the surgery only arm, while 216 to the lymphadenectomy arm.

After a median follow-up of 68.4 months, tumor recurrence was observed in 292 patients (153 in the surgery alone arm, and 139 in the lymphadenectomy arm).

Figure 4 shows the survival curve for these data.




Figure 4: Result of survival analysis for the ovarian cancer trial.

The Number of patients still at risk at each time point for the treated and control groups are shown below the horizontal axis; you can also decide to add the number of events occurred in each interval (e.g. between year 1 and year 2).

The unadjusted hazard ratio for the risk of a first event in the lymphadenectomy group, as compared with the surgery alone group, was 0.76 (95% CI 0.60 to 0.96), meaning that the risk of a recurrence of disease is lower in the patients treated with lymphadenectomy if compared with the control group.

The absolute gain in survival for the treated group is equal to $(1 - 0.76) \times 100 = 24\%$. This difference in survival probabilities of the two compared groups is statistically significant ($P = 0.022$ using the log-rank test).

Median progression-free survival was 22.4 months for patients in the control arm and 27.4 months for patients in the systematic lymphadenectomy arm, confirming the improvement in survival for the patients treated with additional lymphadenectomy.

These results can be obtained using Proc LIFETEST. Proc LIFETEST computes non-parametric estimates of the survival distribution and rank tests for associations of the response variable with other variables. The survival estimates

PhUSE 2008

are computed within defined strata levels and the rank tests are pooled over the strata and are therefore adjusted for strata differences. Additionally, statistics testing of homogeneity of the strata are computed.

```
proc lifetest data=lympho1 plot=(s);
time ttp*censor(0);
strata group;
run;
```

The PLOTS statement displays the survival curve for the data considered. The statement TIME requests the variable for time to progression (*ttp*) and the variable indicating censoring (*censor*); the value between round brackets (0) is the one associated to censored observations, STRATA is the variable for the groups to be compared (*group*). By default the estimates are calculated using the Kaplan-Meier method.

The same code can be used for calculating median of follow-up by changing the censoring indicator from '0' to '1' if '1' indicates the occurrence of the event of interest.

Cox proportional hazards analysis was performed to adjust the treatment comparison for baseline characteristics: results are shown in Figure 5.

When histologic grade and residual disease were taken into account, the adjusted hazard ratio for occurrence of first event was almost unchanged (0.75 [95% CI = 0.59 to 0.94]), confirming that the 25% improvement in survival for patients treated with additional lymphadenectomy depends mostly on treatment received. These results can be obtained using PROC PHREG.

Prognostic factor	PFS	
	HR (95% CI)	<i>P</i>
Treatment arm		
No lymphadenectomy	1.0 (Referent)	.01
Lymphadenectomy	0.75 (0.59–0.94)	
Tumor grade		
1 or 2	1.0 (Referent)	.76
3	1.04 (0.81–1.34)	
Residual tumor		
No	1.0 (Referent)	<.001
Yes	1.65 (1.29–2.12)	

*HR = hazard ratio; CI = confidence interval.

Figure 5: Multivariable Cox proportional hazards analysis for PFS

Proc PHREG procedure performs regression analysis of survival data based on the Cox proportional hazards model. Adopting a partial likelihood function, Cox regression eliminates the unknown baseline hazard and accounts for censoring survival times. It also allows the use time dependent explanatory variable, which is one whose values for any given individual can change over time

```
proc phreg data=lympho2;
model ttp*censor(0)=group grade residual;
run;
```

The statement MODEL requests to specify the dependent (*ttp*, *censor*) and independent (*group*, *grade*, *residual*) variables to fit the regression model; the Hazard ratios are automatically calculated.

EXTRACTING SUMMARY STATISTICS

³ If you want to know which ODS table should be used to access a portion of the output of the SAS procedure you are using, you can simply use the statement ODS TRACE ON before using the procedure and use the statement ODS TRACE OFF when you want to stop ODS tracing. The name of the ODS tables used/generated by the SAS procedure, will be reported in the LOG file.

PhUSE 2008

Survival analysis procedures, such as PROC LIFETEST, are able to generate outputs with a lot of information (the example in figure 2 can increase in size if the study has thousands of patients) and of different types (lifetable, statistical tests). In this case a summary output such as the one in figure 6 may facilitate the life of anyone who needs to interpret the results (the statistician, medical writer, clinician, or the principal investigator). In the summary table in figure 6 the summary statistics describing the 'Time to Progression' across the different study groups are reported:

- nr. of events;
- nr. of censored observations;
- survival estimates at 6, 9 and 12 months (point of clinical interest) with their 95% confidence intervals;
- the 25th, median and the 75th percentiles

This summary table can be generated by extracting the statistics obtained by PROC LIFETEST using the ODS OUTPUT statement. The ODS OUTPUT allows you to convert the output of a SAS procedure into a SAS dataset. For example to extract the Product Limit Estimates we can use the following statements:

```
ods output ProductLimitEstimates=survest;
proc lifetest data=lymphol;
time ttp*censor(0);
strata group;
run;
ods output close;
```

where *ProductLimitEstimates* is the ODS table name and *survest* is the dataset in which the product limit estimates will be saved. Table 1 reports all the ODS tables name³ that can be accessed from PROC LIFETEST. Of course you need to work on the dataset extracted and produce the above table using SAS procedure such as TABULATE (we have for example developed a macro called SURVEST that 'automatically' produce the output in figure 6).

		All Doses (N=13)	Dose Level 1 (N=3)	Dose Level 2 (N=4)	Dose Level 3 (N=6)
Number Progressed/Died/Started New Anti-Tumor Therapy		10 (76.9%)	3 (100.0%)	4 (100.0%)	3 (50.0%)
Number Censored		3 (23.1%)	0 (0.0%)	0 (0.0%)	3 (50.0%)
Kaplan-Meier Estimates of the Proportion of Patient Without PD at	6 Mos	46.7	66.7	0.0	66.7
	9 Mos	24.8	0.0	0.0	66.7
	12 Mos	19.8	0.0	0.0	66.7
25% Percentile of Time to Progression (Months) and its (95% CI)		2.1 (2.0,3.1)	2.0 (2.0,7.6)	1.5 (0.8,3.1)	2.1 (1.8,12.0)
Median Time to Progression (Months) and its (95% CI)		4.8 (2.2,8.3)	6.6 (2.0,7.6)	2.6 (0.8,4.8)	12.0 (2.1,12.0)
75% Percentile of Time to Progression (Months) and its (95% CI)		8.7 (6.3,13.7)	7.6 (2.0,7.6)	4.0 (2.2,4.8)	12.0 (-,-)

Figure 6: Survival Analysis Summary Table

CREATING THE TIME-TO-EVENT ANALYSIS DATABASE

Creating the analysis database (ADB) for time-to-event outcomes may seem an easy task; in theory you just need to create two variables: the event (or censored) variable and the time variable. However, due to the complexity of the clinical trials data collection tools (e.g. CRF), it might not be possible to derive these two variables from only two source variables. In other words it will be uncommon to have a form where the event and the time are directly reported (e.g. event is 0 and time is 30 months), but you need to derive them by accessing different forms (SAS datasets) across different time-schedule (e.g. visits, follow-up), to check whether or not a subject/patient has experienced an event and when. Moreover, you also need to define what has to be considered 'time 0', i.e. either the randomization date or the first diagnosis date, etc. Whilst the event is in most cases something that can be detected without doubts (e.g. the death or the disease progression), the time for censored patients is not.

For example the 'overall survival' is defined as the time from randomization date, or whatever is considered the 'time 0', until the date of death. If the patient is alive at a certain point, the patient is censored at last date it was seen alive (e.g. follow-up date, last phone contact, etc.). Similarly, the 'time-to-progression' endpoint commonly used in oncology, is

PhUSE 2008

calculated from the randomization date till a tumor progression is observed. In the event that a patient did not experience tumor progression, the patient will be censored either at the last contact date, or the last date when a formal tumor-assessment was performed, which may or may not be the same as the last date the patient was seen alive. As this specific situation may occur with other study-endpoints in other therapeutic areas, we suggest to carefully define the algorithm to be used by clearly detailing the situation like the one just described. In appendix, you can find an example of 'ADB derivation' flow-chart.

ODS Tables Name	Contents
CensoredSummary	Number of events and censored observations
HomStats	Rank statistics for testing strata homogeneity
HomTests	Tests for strata homogeneity
LifetableEstimates	Lifetable survival estimates
LogForStepSeq	Forward stepwise sequence for the log-rank statistics for association
LogHomCov	Covariance matrix for the log-rank statistics for strata homogeneity
LogTestCov	Covariance matrix for log-rank statistics for association
LogUniChisq	Univariate chi-squares for log-rank statistics for association
Means	Mean and Standard Error of survival time
ProductLimitEstimates	Product-limit survival estimates
Quartiles	Quartiles of the survival estimates
WilForStepSeq	Forward stepwise sequence for the log-rank statistics for association
WilHomCov	Covariance matrix for the Wilcoxon statistics for strata homogeneity
WilTestCov	Covariance matrix for the log-rank statistics for association
WilUniChisq	Univariate chi-squares for the Wilcoxon statistic for association

Table 1: ODS Tables Produced by PROC LIFETEST

ODS STATISTICAL GRAPHICS IN SAS VERSION 9.2

ODS Graphics is a new feature in SAS Version 9.2 (it was experimental in SAS Version 9.1), from SAS products of SAS/STAT, Base SAS, SAS/ETS, and others, which allows the users to create statistical graphics automatically. Both PROC LIFETEST and PROC PHREG support ODS Graphics. For example PROC LIFETEST provides eight type of ODS GRAPHICS. Figure 7 shows the default "Estimated Survival Curve" (ODS Graph Name "Survival") generated using the following SAS code:

```
ods graphics on;
proc lifetest data=surv plot=(s);
time ttp*censor(0);
strata group;
run;
quit;
```

Prior to version 9.2 and ODS GRAPHICS, a similar output could be produced using the following code:

```
proc lifetest data=surv plot=(s)graphics;
time ttp*censor(0);
strata group;
symbol1 line=1 color=blue;
symbol2 line=2 color=red;
run;
```

PhUSE 2008




Figure 7: Graphical Survival Plot with ODS Graphics

META-ANALYSIS OF TIME-TO-EVENT DATA

Meta-analyses (MA) that combine the results of related randomised controlled trials have become a common and widely accepted tool in the evaluation of health care interventions. Regulatory agencies such as the FDA have started considering it as key supporting evidence. Systematic reviews use explicit, objective and prospectively-defined methods to collect, critically evaluate and synthesize studies, making them less biased and more reproducible than traditional reviews. A systematic review may, or may not, include a meta-analysis (MA): a statistical pooling of the results from individual studies to obtain a single overall estimate of treatment effect. Today there are well validated method to perform also meta-analysis of time-to-event data [13] [14].

ADDITIONAL TECHNICAL REFERENCES

Apart the methodological references you can find in the 'References' section, we want to suggest to have a look at the following articles presented at various SAS users congress:

- Smith T. Smith B. Kaplan Meier and Cox Proportional Hazard Modelling: Hands on Survival Analysis. WUSS. 2004
- Yeh S. ODS and Statistical Graphics. Paper PR02. NESUG17. 2004
- Zhou JC. Enhancement of Survival Graphs. Paper 239-28. SUGI28
- Xu W. Wang WWB %KMPLLOT – An enhancement of SAS PROC LIFETEST. Paper SP06. PharmaSUG. 2005
- Smoak C. Survival (Kaplan-Meier) Curves made Easy. Paper TT07. PharmaSUG. 2006
- Gharibvand L. Fernandez G. Advanced Statistical and Graphics features of SAS PHREG. Paper 375. SAS Global Forum. 2008

CONCLUSIONS

Survival analysis is a method of analyzing time to event data when such data are exposed to a censoring process, whereby subjects finish a clinical trial without having experienced an event. Other statistical methods cannot adequately take into account the censoring process since they do not adjust for the 'risk' population decreasing as the subject leaves the trial. Statistical procedures such as Proc LIFETEST and Proc PHREG are the tools with which the statistician/programmer has to work with within the SAS system. Statistical programmers require a fundamental understanding of the methodology so that they are able to understand, and eventually check, the results that they are without doubts responsible for.

REFERENCES

1. Kaplan EL. Meier P. Nonparametric estimation from incomplete observations. J Am Statistics Association 1958; 53:457-81
2. Bewick V., Cheek L., Ball J. Statistics review 12: survival analysis. Crit Care 2004 Oct;8(5):389-94. Epub 2004 Sep 6.
3. Douglas G Altman, J Martin Bland. Statistics Notes: Time to event (survival) data. BMJ 1998; 317: 468-469
4. J Martin Bland , Douglas G Altman. Statistics Notes: Survival probabilities (the Kaplan-Meier method). BMJ 1998;

PhUSE 2008

- 317; 1572-1580
5. Parmar M. and Machin D. Survival Analysis – A Practical Approach. 1995, John Wiley & Sons, Chichester, England
 6. TG Clark, MJ Bradburn, SB Love and DG Altman. Survival Analysis Part I: Basic concepts and first analyses. British Journal of Cancer (2003) 89, 232-238
 7. Kleinbaum DG: Survival Analysis: A self-learning Text. Statistics in the Health Sciences. Springer 1997
 8. J Martin Bland , Douglas G Altman. The logrank test. BMJ 2004; 328; 1073
 9. Cox DR. Regression models and life tables (with discussion). Journal of the Royal Statistics Society, B34, 187-220
 10. Cox DR. Oakes D. Analysis of Survival Data. Chapman & Hall, London. 1984
 11. Spotswood L. Spruance, Julia E. Reid, Michael Grace, Matthew Samore. Hazard Ratio in Clinical Trials. Antitumoral Agents and Chemotherapy Aug. 2004; 2787-2792
 12. Benedetti Panici P. et al. Systematic aortic and pelvic lymphadenectomy versus resection of bulky nodes only in optimally debulked advanced ovarian cancer: a randomized clinical trial. Journal of the National Cancer Institute, Vol. 97, No. 8
 13. Parmar KB, Torri V, Stewart L. Extracting summary statistics to perform meta-analyses of the published literature for survival endpoints. Stat Med. 1998; 17: 2815-2834
 14. Tinazzi A., Tierney J. A gentle introduction to meta-analysis. PhUSE 2007, Paper TU03

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Angelo Tinazzi

SENDO Tech S.r.l.

Medical Informatics and Biometry Unit

Via Visconti di Modrone 12

Milan / 20122

Work Phone: 0039 02 76420426

Fax: 0039 02 76017484

Email: tinazzia@sendo-org.it

Web: <http://www.sendofoundation.org>

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

APPENDIX: EXAMPLE OF A FLOW-CHART OF A COMPLEX TIME TO EVENT ADB CREATION

The following flow-chart summarize the programming steps to be followed to derive the event and time variables for the time-to-progression(PD) endpoint.

