

Forest Plots in Survival Subgroup Analyses

Julia Benzing, Chrestos Concept GmbH & Co.KG, Essen, Germany
Michael Brockmann, Chrestos Concept GmbH & Co.KG, Essen, Germany

ABSTRACT

In this paper we demonstrate how forest plots can be used in a comparative subgroup analysis. For different subgroup populations we compute median overall survival times, survival rates at different time points, hazard ratios, as well as corresponding confidence intervals based on a Cox proportional hazards model, which can be modified by stratification factors. The calculations are performed by a SAS macro which we developed for the purpose of comparing such efficacy parameters and which is based on the procedures PROC LIFETEST and PROC PHREG in SAS 9.4. The macro automatically generates a forest plot that nicely displays and compares those statistical parameters among different subgroup populations. It further provides appropriate tables with additional survival analysis information such as number of patients at risk and p-values. We will show some results based on dummy data constructed in the style of typical time to event data from oncology.

INTRODUCTION

In the survival analysis of oncological trials, a comparison of median survival times and survival rates at fixed times between different treatment arms is of paramount interest. One objective is the survival time, which is the time interval from first exposure to treatment until the first event occurs. In this context, an event is usually specified by the definition of so-called “endpoints”. In this paper we are examining two different endpoints, “progression-free survival” (PFS) and “overall survival” (OS), which are both typical in oncological trials. Another objective is the survival rate at a fixed time, which is the probability of not experiencing an event within a certain time after first exposure to treatment, e.g., in the case of OS, to be still alive at that time.

Unfortunately, data sets containing information like disease progression or death are almost always incomplete and special techniques are required to analyze the data. For instance, parameters like the median survival time and the survival rate at a fixed time can be calculated by means of the so-called “Kaplan-Meier estimator” [1]. This is also known as the “product limit estimator”, which does not need any parameters as input. For each treatment arm the median survival times or the survival rates at a fixed time can be easily read out from the Kaplan-Meier curve [2] and can subsequently be compared to each other. The Kaplan-Meier curve is simply the graphical representation of the survival rate as function of time. It monotonically decreases. If one of the Kaplan-Meier curves is well above the other at all times, this strongly indicates that one treatment is more effective than the other.

On the other hand, if two Kaplan-Meier curves are close together or if they intersect at one or more time points, such an interpretation becomes questionable. Caution has to be taken if we have to compare survival rates at a certain time point. The result might not be statistically significant. In order to overcome this problem and to decide whether one treatment is significantly better than the other, we have several methods at hand. One of them uses the so-called “hazard ratio” [3]. It is defined as the ratio of hazard rates in two treatment arms. For instance, if patients receiving the first treatment die at twice the rate per unit time as the control group, the hazard ratio will be two. The corresponding confidence interval at a given significance level α is defined as the interval in which a repeatedly measured hazard ratio would lie with a probability of $1-\alpha$. If the confidence interval does not include unity, the hazard ratio is considered “statistically significant (at level α)”. We will explicitly define the hazard ratio and its confidence interval in the next section.

In the sample data of this paper we are examining two treatment arms, which we allude to in the following as “Study drug” and “Placebo”. This divides the entire population into two treatment groups. Their survival rates can be compared to each other by using the hazard ratio. Aside from treatment information there might be more categories that divide the entire population into subgroups, e.g. “Gender” or “Age group”. More details about the data structure and, in particular, how occurrences of events and event times are stored in this “time-to-event” data will be presented in a later section.

Eventually, we will compare hazard ratios and their confidence intervals among different subgroups. For this purpose a so-called “forest plot” is proven to be useful [4]. It displays the hazard ratio as a dot and its confidence interval by bars like error bars. If plotted vertically, a collection of such lines resembles a bunch of trees and is therefore called “forest plot” [5]. We have developed a macro in SAS[®] 9.4 that is capable of creating this forest plot by solely using the time-to-event data as input, provided that the structure of data follows common standards (i.e. CDISC standard [6]). We will present a couple of interesting features and more details of the macro in the main section of this paper.

In conclusion, we will summarize the results. Furthermore, as a short overview, we will present some possible features that might be added to the macro in future applications.

THEORETICAL BACKGROUND

Relative risks are applicable in many contexts of clinical trials. Their interpretation is easy. One can use them, for instance, to compute the quotient of likelihoods of death under treatment A and under treatment B. The use of relative risks in the context of survival times, on the other hand, might not be appropriate because there might be different censorships. Thus, there are no exact time intervals. This is the reason why a measure with a time component is needed.

FROM THE LIFE TABLE APPROACH TO THE HAZARD RATIO

The first step is building time intervals:

$$[t^{(0)}, t^{(1)}], [t^{(1)}, t^{(2)}], [t^{(2)}, t^{(3)}], \dots$$

This is called the life-table approach. The relative risk of the time interval $[t^{(p)}, t^{(p+1)}]$, $p = 0, 1, 2, \dots$, is defined as the quotient of the conditional probabilities of an event under treatment A and under treatment B, given a patient has not yet experienced an event at time point $t^{(p)}$. Assuming that the relative risks are independent of the time interval, a definition of a global measure is possible. This allows an estimate of the relative risks of variable time intervals [7].

The limit of infinitesimally small time intervals of the conditional probability of an event is called hazard function:

$$\lambda(t) = \lim_{h \rightarrow 0} \frac{P(t < T \leq t+h \mid T > t)}{h}.$$

In this formula, T serves as the patient's time of event. The assumption of proportional hazards is that the two hazard functions of different treatments are proportional to each other. This implies that their quotient is the same at every time point – along with the assumption we made for the relative risks. The hazard ratio, in the following called HR , is given by

$$HR = \frac{\lambda_B(t)}{\lambda_A(t)} = \text{const.}$$

This means that the effects of treatments A and B, e.g. in a clinical trial, are the same throughout the complete observation period [7].

COX PROPORTIONAL HAZARDS MODEL

The Cox proportional hazards (PH) model is a regression model that can be used to analyze survival data. The model can be formulated as

$$\lambda(t, X) = \lambda_0(t) \exp(\sum_{i=1}^p \beta_i X_i),$$

where $X = (X_1, X_2, \dots, X_p)$ is a collection of predictor variables, e.g. treatment, height, or age. The baseline hazard function λ_0 is an unspecified function of time but is independent of the X 's. The part $\exp(\sum_{i=1}^p \beta_i X_i)$ is the linear sum of the p explanatory variables X , which are time-independent. If they depend on time, the assumption of proportional hazards is not fulfilled anymore and an extended version of the Cox model should be used [8].

Since λ_0 is an unspecified function, the Cox PH model is a semi-parametric model. The correct parametric model is generally not known. The Cox model is a “robust” model, which means that the results of the model closely approximate the results of the correct parametric model. That is why the Cox model is called a “safe option” [8].

The parameters $\beta = (\beta_1, \beta_2, \dots, \beta_p)$ can be estimated by the maximum likelihood method. Within the Cox model the likelihood function $L(\beta)$ is not a “complete” but a “partial” likelihood function. It examines probabilities of patients failing and does not take into account the probabilities of patients being censored. The total likelihood L is the product of likelihoods L_j of failing at time point j , given the patient has not experienced an event before:

$$L = \prod_{j=1}^k L_j .$$

In this formula, k represents the number of failure times.

The set of individuals at risk for the current time point is called risk set $R(t^{(m)})$. Therefore, the likelihood only counts patients who fail. The survival time information prior to censorship is used for censored individuals. Therefore, a patient who is censored after time point $t^{(m)}$ is in $R(t^{(m)})$. The maximum likelihood can be estimated by using iteration [8].

STRATIFIED COX MODEL

If a predictor does not satisfy the assumption of proportional hazards, the stratified Cox model should be applied. It is a modification of the Cox PH model that includes the predictors fulfilling that assumption and does not include all other predictors. The Wald test, for instance, helps to find out which variables should be excluded from the model. One example could be the variable “Gender”. Treatment often has different effects on females and males. It thus makes sense to use this variable as a stratification factor in the model.

The hazard function of a stratified Cox model is defined as follows

$$h_g(t, X) = h_{0g}(t) \exp\left(\sum_{i=1}^p \beta_i X_i\right) ,$$

with p being the number of variables fulfilling the assumption, and g being a subscript indicating the stratum. Consequently, every characteristic of the stratified variable has a different hazard baseline function. For the variable “Gender”, for instance, we have $g = F, M$.

As in the unstratified model, the (partial) maximum likelihood method can be used to estimate the β 's. In the stratified model the likelihood is the product of the likelihoods of every stratum [8].

WALD TEST

The Wald test can be used to test whether a regression coefficient is significantly different from zero. In order to calculate the test statistic, the estimated parameter is divided by its standard error:

$$W = \hat{\beta}_i / SE(\hat{\beta}_i) .$$

In the hypothesis that $\beta_j = 0$, the test statistic W has a standard normal distribution. The $100(1 - \alpha)\%$ confidence interval of the hazard ratio can be computed by means of

$$\exp\left(\hat{\beta}_i \pm u_{1 - \frac{\alpha}{2}} SE(\hat{\beta}_i)\right) ,$$

whereby $u_{1 - \frac{\alpha}{2}}$ is the $\left(1 - \frac{\alpha}{2}\right)$ – quantile of a standard normal distribution [7].

SURVIVAL DATA

The sample data in this paper consists of 210 patients, each of which has a unique identification number (SUBJIDN). Each patient received one of two different treatments: "Study Drug" or "Placebo", which are encoded in the variable "Treatment" (TRT01PN). Additionally, there are three different categories that divide the entire population into subgroups: "Gender", "Age Group", and "ECOG PS at baseline" (ECOG = "Eastern Cooperative Oncology Group"; PS = "Performance Status").

In our example of time-to-event data (ADTTE) displayed below, we examine two different endpoints: "Progression-Free Survival" (PFS) and "Overall Survival" (OS). According to the CDISC standard of time-to-event data [6], endpoints can be distinguished by the value of the "Parameter code" (i.e. PARAMCD="PFS" and PARAMCD="DEATH", respectively; see Table 1 below).

In an ideal trial the time to an endpoint is recorded for every patient. Unfortunately, time-to-event data is almost always incomplete. For example, patients might withdraw from the trial or die before any progression is observed. For this reason, a censoring variable is added to the time-to-event data. Its value is 0 for an event actually happening and 1 for every lost patient. Hence, the data set has the following structure (see Table 1): For each patient there are two observations, one for each endpoint. The variables containing the full survival information are the "Censor variable" (CNSR $\in \{0,1\}$) and time from first treatment date to the event/censoring date, depending on the value of CNSR. This time interval is stored in the variables "Analysis value" (AVAL) and "Analysis value unit" (AVALU).

ADTTE dummy data (illustrative subset)

Obs	Subject identifier	Gender	Age group	ECOG PS at baseline	Treatment	Parameter code	Censor variable	Analysis value	Analysis value unit
1	1401920001	Female	>= 85	0-1	Study drug	PFS	0	301	Days
2	1401920001	Female	>= 85	0-1	Study drug	DEATH	1	315	Days
3	1401920003	Male	65-74	0-1	Placebo	PFS	1	126	Days
4	1401920003	Male	65-74	0-1	Placebo	DEATH	1	126	Days
5	1401920107	Male	75-84	2-3	Study drug	PFS	0	61	Days
6	1401920107	Male	75-84	2-3	Study drug	DEATH	1	61	Days
7	1401920113	Female	65-74	2-3	Study drug	PFS	1	1	Days
8	1401920113	Female	65-74	2-3	Study drug	DEATH	1	1	Days
9	1401920196	Female	< 65	4	Study drug	PFS	0	51	Days
10	1401920196	Female	< 65	4	Study drug	DEATH	0	94	Days
11	1401920204	Male	75-84	4	Placebo	PFS	0	19	Days
12	1401920204	Male	75-84	4	Placebo	DEATH	0	19	Days
13	1401920210	Female	75-84	.	Placebo	PFS	0	23	Days
14	1401920210	Female	75-84	.	Placebo	DEATH	1	23	Days

Table 1: Only part of the time-to-event data (ADTTE) is listed. The original data consists of 210 patients. The seven patients selected here show possible values of subgroup variables "Gender", "Age group", and "ECOG PS at baseline" as well as the treatment information for each patient ("Treatment"). It further demonstrates possible combinations of endpoints "PFS" and "DEATH" ("Parameter code") and the "Censor variable" using values 0 or 1. The time to event and its unit are stored in the variables "Analysis value" and "Analysis value unit".

Note that subgroup or treatment data are not always complete. In our example, the patient with SUBJIDN=1401920210 (see last two rows of Table 1) has a missing value of "ECOG PS at baseline". One way to handle missing data is to follow the rules in the statistical analysis plan (SAP). In the survival analysis in the next

section, patients having only partial subgroup information are removed from the whole analysis. Of course, different rules could be easily implemented in the macro presented in the next section.

SAS MACRO CREATING FOREST PLOTS

In this section we present some details and interesting features of the macro that we developed in SAS® 9.4 [9], in order to analyze survival data for various subgroups. The macro is capable of performing a comparative subgroup analysis of time-to-event data (ADTTE), provided the data has the structure as described in the previous section. It computes median survival times and survival rates at fixed times by using PROC LIFETEST as well as hazard ratios and corresponding confidence intervals by using PROC PHREG. The latter is based on a Cox proportional hazards model that can be stratified by stratification factors. The output of the macro is a forest plot that nicely displays and compares the survival data among various subgroups. The subgroups are fixed by the categorical variables that are handed over to the macro.

OVERVIEW OF MACRO VARIABLES

The macro is called %m_forest_plot. It provides a lot of input variables that can be used to select the data and to fix the layout of the result. The most important macro variables, which we will use later in this section, are listed in the following lines of code:

```
%m_forest_plot(
    indat          = adtte      /*@type DATA*/
    , endpoint      =          /*@type TEXT (should be value of PARAMCD)*/
    , subgroup_varlist =       /*@type VARIABLES(#data)*/
    , subgroup_labels =       /*@type TEXT (labels separated by '~') */
    , time_unit_display =     /*@type ONE( |Days|Weeks|Months|Years)*/
    , time_format    =       /*@type FORMAT (to display NUMBERS)*/
    , strat_factors   =       /*@type VARIABLES(#data)*/
    , strat_fact_ftntxt =     /*@type TEXT*/
    , ci_level        = 95     /*@type NUMBER (0 < ci_level < 100)*/

    , show_events_cnsr = yes   /*@type ONE( |YES|NO)*/
    , show_median_surv = yes   /*@type ONE( |YES|NO)*/
    , show_surv_rates  = no    /*@type ONE( |YES|NO)*/
    , time_points      =       /*@type NUMBERS (seperated by blanks)*/

    , show_extra_ftn   = yes   /*@type ONE( |YES|NO)*/
    , show_subj_list    = no    /*@type ONE( |YES|NO)*/
    , show_ftn_median   = yes   /*@type ONE( |YES|NO)*/
    , show_ftns_hr_ci   = yes   /*@type ONE( |YES|NO)*/
    , subj_display      = subjidn /*@type VARIABLE(#data)*/
    , footnote1         =       /*@type TEXT*/
    , footnote2         =       /*@type TEXT*/
    , footnote3         =       /*@type TEXT*/
    , footnote4         =       /*@type TEXT*/
    , footnote5         =       /*@type TEXT*/

    , print_data_list   = yes   /*@type ONE( |YES|NO)*/
    , outdat            =       /*@type DATA*/
    , forest_plot_ny    = yes   /*@type ONE( |YES|NO)*/
    , forest_plot_file  =       /*@type TEXT*/
    , overall_ny        = yes   /*@type ONE( |YES|NO)*/

) / DES = 'Creating a forest plot of efficacy data by subgroup';
```

The first part of the macro code (about one third) is dedicated to consistency checks. In case of an inconsistency, the macro informs the developer by writing error messages, warnings, or notes to the log. By default the macro takes the data set ADTTE from the library WORK as input. It is of course possible to use a different input data set by assigning its name to the macro variable indat. If the data set assigned to indat does not exist, the macro aborts and prints an error message to the log. Furthermore, it checks if there is a variable PARAMCD with value "&endpoint" in &indat. If this is not the case, it again aborts and prints an error message to the log. So it does if one of the subgroup variables listed in subgroup_varlist is not part of &indat, and so on and so forth.

CALCULATING SURVIVAL DATA

In the second part of the code the survival data is calculated. The following PROC PHREG step demonstrates how the hazard ratio and its confidence interval are computed by the macro. The category `&by` is given by one of the subgroup variables listed in the macro string `&subgroup_varlist`.

```
/* (Un-)stratified Cox proportional hazards model */
ODS OUTPUT ParameterEstimates = hazard_ratio

/* Original PROC PHREG step from the macro... */
PROC PHREG data=&indat.1 alpha=&alpha.;
    CLASS &treat_var.;
    MODEL &time.*&censor_var.(&censor_value.)=&treat_var. / rl;
    %IF %length(&strat_factors.) %THEN STRATA &strat_factors.;;
    BY &by.;
RUN;

/* ...and how it looks like after resolving all macro variables: */
PROC PHREG data=adttel alpha=0.05;
    CLASS trt01pn;
    MODEL aval*cnsr(1)= trt01pn / rl;
    [e.g. STRATA sex agegr01n;] /* Depending on length of &strat_factors. */
    BY ecoggrp;
RUN;
```

The first block is the PROC PHREG step as it originally appears in the macro code, whereas the second block shows how it appears if all macro variables are resolved. In this example, `&by` is given by the subgroup variable ECOGRPN. Resolution of most of the other macro variables is straightforward, provided that the data `&indat.1` has the common structure of ADTTE as described in the previous section. The confidence level `alpha=%eval(1-&ci_level./100)` is determined by the macro variable `ci_level`. Default values are `&ci_level=95` and thus `&alpha=0.05`. The suffix 1 at `&indat.1` indicates that little preparation of the input data set `&indat` is necessary before we can use it in the PROC PHREG step. For example, if the unit of time variable AVAL is AVALU="Weeks", rather than AVALU="Days", the macro re-calculates AVAL (by multiplying with 7), in such that "Days" can be used as a unit. If we want to display median survival times in a different fashion, we can realize this by simply assigning one of the values `weeks`, `months`, or `years` to the macro variable `time_unit_display`. The macro then transforms all values of AVAL to fit into the desired unit.

FOREST PLOTS AS OUTPUT OF THE MACRO

Once completed, the computation of the survival parameters (e.g. number of events, number of censored patients, median survival times and so on), the survival data will be transcribed to a data set. After further processing, this data will be used to create the forest plot by means of PROC SGPLOT. Before applying this PROC SGPLOT step, all footnotes and titles have to be set. For this purpose there are several switches in the list of macro variables, all of which begin with "show_" (see above). They enable us to decide which footnotes should appear below the forest plot. Some examples are shown in the figures below. For instance, calling the macro as in the following yields Fig. 1:

```
%m_forest_plot(
    endpoint          = pfs
    , subgroup_varlist = sex agegr01n ecoggrp
    , subgroup_labels  = Gender~Age Group~ECOG PS
    , time_unit_display = Days
    , footnotel        = %nrstr(Footnote example: The macro can handle up to
                                five additional footnotes, to be specified by
                                macro variables footnotel-footnote5.)
    , forest_plot_file = Forest_Plot_PFS_95
);
```




Figure 1: The figure shows the hazard ratios of all subgroups in the four categories “Overall”, “Gender”, “Age group”, and “ECOG PS”, together with their confidence intervals, illustrated by horizontal bars in a logarithmic scale. The length of a blue thick bar is proportional to the size of the corresponding subgroup population. In the table next to the forest plot additional survival information like median survival times of the two different treatment arms are illustrated. The red vertical line marks a hazard ratio of one.

The last footnote appears because there is one patient in the input data who has a missing value of “ECOG PS at baseline” (see Table 1 in the previous section). This is how the footnote can appear in the final report. During development of the figure, on the other hand, it would be nice to know which patients have missing values. To this end, there is a macro variable `show_subj_list`, which can be switched from `no` (default) to `yes`. In the next example we also switch from confidence level `alpha=95` to `alpha=80` and from endpoint “Progression-Free Survival” to “Overall Survival” by changing `endpoint` from `pfs` to `death`. Both cases illustrate how title and subtitle automatically change. Note that a few typical endpoints of survival analysis are already implemented in the macro. More exotic endpoints can easily be added. The following macro call generates Fig. 2:

```
%m_forest_plot(
  endpoint                = death
  , subgroup_varlist      = sex agegr01n ecoggrp
  , subgroup_labels       = Gender~Age Group~ECOG PS
  , time_unit_display     = weeks
  , ci_level              = 80
  , show_subj_list        = yes
  , footnote1             = %nrstr(Footnote example: The macro can handle up to
                                five additional footnotes, to be specified by
                                macro variables footnote1-footnote5.)
  , forest_plot_file      = Forest_Plot_PFS_95
);
```




Figure 2: This figure is similar to Fig. 1. One difference is the choice of the endpoint, which is “Overall survival” instead of “Progression-free survival”, another one is the change of the time unit of the median survival times. Furthermore, there is an additional footnote about missing survival data, and the last footnote has slightly changed due to the extra option `show_subj_list=yes` in the macro.

Another difference to the previous example (cf. Fig. 1) is that we changed the value of the macro variable `time_unit_display` from `days` (default) to `weeks`. This effects how median survival times are displayed in the table next to the forest plot. Depending on its value (`days`, `weeks`, `months`, or `years`) an appropriate format is automatically chosen by the macro. If another format is desired, one can overwrite it by explicitly assigning a format to the macro variable `time_format`. For example, the choice `time_format = 5.1` displays survival times in the SAS format 5.1.

So far, we have not considered any stratification yet. If one suspects that one or more variables do not fulfill the assumption of proportional hazards, stratification of the Cox model by that/those variable(s) is required. Assigning those variables to the macro string `strat_factors` can accomplish this. The following macro call yields Fig. 3:

```
%m_forest_plot(
  endpoint          = pfs
  , subgroup_varlist = sex agegr01n ecoggrp
  , subgroup_labels  = Gender~Age Group~ECOG PS
  , time_unit_display = years
  , ci_level         = 95
  , strat_factors    = sex agegr01n
  , strat_fact_ftntxt = %nrstr(Gender and Age Group)
  , footnote1        = %nrstr(Footnote example: The macro can handle up to five
  additional footnotes, to be specified by macro variables footnote1-footnote5.)
  , overall_ny       = no
  , forest_plot_file  = Forest_Plot_PFS_95_strata
);
```




Figure 3: In this forest plot the hazard ratios and their confidence intervals were estimated based on a stratified Cox proportional hazards model, as mentioned in one of the footnotes. Furthermore, the category “Overall” has been removed and the time unit has been changed again to years.

Note that the footnote citing the Cox model now emphasises the stratification and lists the stratification factors, according to the value of `strat_fact_ftntxt`. Another difference to the previous figure is that the row “Overall” is removed due to the option `overall_ny=no`.

The layout changes more drastically if we change the content of columns in the table next to the forest plot. For instance, it is possible to add additional columns by changing the variables `show_events_cnsr` and `show_surv_rates` from their default values `yes` and `no` and `yes`, respectively. The option `show_surv_rates=yes` requires to specify time points for which the survival rates should be computed. This can be done by listing them in the macro string `time_points`. The unit of `time_points` will be automatically set to the same unit as defined in `time_unit_display`. Calling the macro in the following way produces Fig. 4:

```
%m_forest_plot(
  endpoint          = death
  , subgroup_varlist = sex agegr01n ecoggrp
  , subgroup_labels  = Gender~Age Group~ECOG PS
  , time_unit_display = months
  , ci_level         = 80
  , footnote1        = %nrstr(Footnote example: The macro can handle up to five
additional footnotes, to be specified by macro variables footnote1-footnote5.)
  , forest_plot_file = Forest_Plot_OS_95_surv_rates
  , show_events_cnsr = no
  , show_surv_rates  = yes
  , time_points      = 3 6
);
```




Figure 4: The forest plot is the same as in Fig. 2. The main difference is that the table next to it has completely changed. Now, it displays survival rates at selected time points for both treatment arms. The footnote citing missing values has been adapted accordingly.

In this last example of the section we again see that the footnote citing missing survival data is automatically set below the plot. We already implemented an additional switch in the macro code that can be used if this footnote should not appear, e.g., in a final report.

CONCLUSIONS

In clinical trials, especially in oncology, survival data like overall survival or time to event data are the main objectives to measure the efficacy of a new drug. It is often desired to present efficacy data in a comparative way, for instance, median survival times or hazard ratios (treatment vs. placebo). For this reason we developed a SAS macro that is capable of calculating hazard ratios of various subgroups as well as of the whole population (optional). The hazard ratios, together with their confidence intervals, are depicted in a forest plot on a logarithmic scale. Which subgroups are displayed depends on the categories that are specified in the corresponding macro string (`subgroup_varlist`). Next to the plot a table with various survival data is shown. It contains, for instance, the number of patients with an event, the number of censored patients as well as statistical parameters like median survival times or survival rates at selected time points.

Furthermore, if the Cox model is stratified, this will be cited in a footnote, which is then accordingly adapted by the macro. We would like to stress that this kind of dynamical output generation is implemented in various places in the macro code. For example, if certain statistical parameters are not estimable due to censored data, an automatically generated footnote appears.

Extensions to the macro are of course possible. Up to now, the macro uses "Treatment" as an effect variable in

the Cox model (see the corresponding formula in the theoretical background section and the MODEL option in the PHREG procedure in the main section of this paper). In our example, the treatment variable takes just two different values, representing the two different treatment arms. However, in some clinical trials there are more than two treatment arms. The macro would then have to be extended accordingly. Alternatively, in order to look for an effect of a certain category, a multivariate analysis might be helpful. This option would then have to be added to the macro as well.

REFERENCES

- [1] Kaplan, E.L. and Meier, P. (1958). *Non-parametric estimation from incomplete observations*. JASA **53** (282): 457–481.
- [2] Rich JT, Neely JG, Paniello RC, Voelker CC, Nussenbaum B, and Wang EW (2010). *A practical guide to understanding Kaplan–Meier curves*. Otolaryngol Head Neck Surg. **143** (3): 331–6.
- [3] Breslow, N.E. (1975). *Analysis of Survival Data under the Proportional Hazards Model*. Rev. Int. Stat. **43** (1): 45–57; Reid JE, Grace M, and Samore M (August 2004). *Hazard Ratio on Clinical Trials*. AAC **48** (8): 2787–2792.
- [4] Hedges, L. V. and Olkin, I. (1985). *Statistical Methods for Meta-Analysis*. Orlando: Academic Press.
- [5] Lewis, S. and Clarke, M. (June 2001). *Forest plots: trying to see the wood and the trees*. BMJ **322** (7300): 1479–1480
- [6] Clinical data interchange standards consortium (CDISC), <http://www.cdisc.org/>; *Study Data Tabulation Model Implementation Guide (SDTMIG) v3.2*, <https://www.cdisc.org/standards/foundational/sdtmig>.
- [7] Schumacher, M. and Schulgen, G. (2008). *Methodik klinischer Studien – Methodische Grundlagen der Planung, Durchführung und Auswertung*. Springer: Berlin.
- [8] Kleinbaum, D.G. and Klein, M (2012). *Survival Analysis – A Self-Learning Text*. Springer: New York.
- [9] The output/code for this paper was generated using SAS software. Copyright © 2018 SAS Institute Inc. SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc., Cary, NC, USA.

CONTACT INFORMATION

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

Author Name	Michael Brockmann
Company	CHRESTOS Concept GmbH & Co.KG
Address	Giradetstraße 1-5
City / Postcode	45131 Essen
Work Phone:	+49 151 70239205
Email:	michael.brockmann@chrestos.de
Web:	www.chrestos.de

Author Name	Julia Benzing
Company	CHRESTOS Concept GmbH & Co.KG
Address	Giradetstraße 1-5
City / Postcode	45131 Essen
Work Phone:	+49 160 97844211
Email:	julia.benzing@chrestos.de
Web:	www.chrestos.de

Brand and product names are trademarks of their respective companies.