

Survival analysis

Patrick Breheny

December 8, 2016

Today's lab is about survival analysis. Obviously, survival analysis is a big subject and we're just scratching the surface, but all biostatisticians should be familiar with what survival data looks like and how it is organized, and be able to construct a Kaplan-Meier curve and carry out some basic inference concerning it.

1 Kaplan-Meier curves

1.1 Basic curves

Survival analysis functions are not loaded by default in R; to access them you have to load the `survival` package. The package is very rich with features and can do many things, but we're just going to use a few of its most common functions. First, let's take a look at the data and how it is organized (using the same anemia data we looked at in class):

```
> require(survival)
> anemia <- read.delim("http://myweb.uiowa.edu/pbreheny/data/anemia.txt")
> head(anemia)
```

	Trt	Time_gvhd	Status_gvhd	Time	Status
1	MTX	9	1	30	1
2	MTX	11	1	44	1
3	MTX	12	1	104	1
4	MTX	20	1	106	1
5	MTX	20	1	181	1
6	MTX	25	1	329	0

As we discussed in class, survival outcomes consist of two components: the time on study and the event/censoring indicator. In the data, these appear as two distinct columns (there are actually four “outcome” columns here, two for GVHD and two for mortality). To analyze the data, we need to “bundle” these columns together into what the `survival` package calls a `Surv` object:

```
> S.gvhd <- with(anemia, Surv(Time_gvhd, Status_gvhd))
> S.mort <- with(anemia, Surv(Time, Status))
> head(S.mort)
```

[1] 30 44 104 106 181 329+

Note that the censored observations appear with a little + after them to indicate that the true time-to-event is over 329 days. Now that we've created the survival outcomes, we can analyze them using `survfit`, which works like the other R functions we've seen in class:

```

> Trt <- anemia$Trt
> fit <- survfit(S.mort ~ Trt)
> fit

Call: survfit(formula = S.mort ~ Trt)

           n events median 0.95LCL 0.95UCL
Trt=MTX    24     9   719    371    NA
Trt=MTX+CSP 22     4    NA     NA     NA

```


This tells us the sample size, the number of events (here, deaths) in each group, and tries to estimate the median survival time and give us a confidence interval for it.

What's going on with the NA's? The reason for them will become clear once we plot the Kaplan-Meier curves. To do so, we can simply call `plot(fit)`; here is that call, with a few options turned on and off so you can see what they do:

```

> plot(fit, xlab="Days on study", ylab="Survival")


```




```

> col <- c("#FF4E37FF", "#008DFFFF")
> plot(fit, xlab="Days on study", ylab="Survival", mark.time=FALSE,
+       col=col, lwd=2, las=1, bty="n")
> text(1000, 0.9, "MTX+CSP")
> text(800, 0.4, "MTX")

```



```
> plot(fit, xlab="Days on study", ylab="Survival", mark.time=FALSE, conf.int=TRUE,
+       col=col, las=1)
> legend("bottomleft", legend=unique(Trt), col=col, lty=1)
```

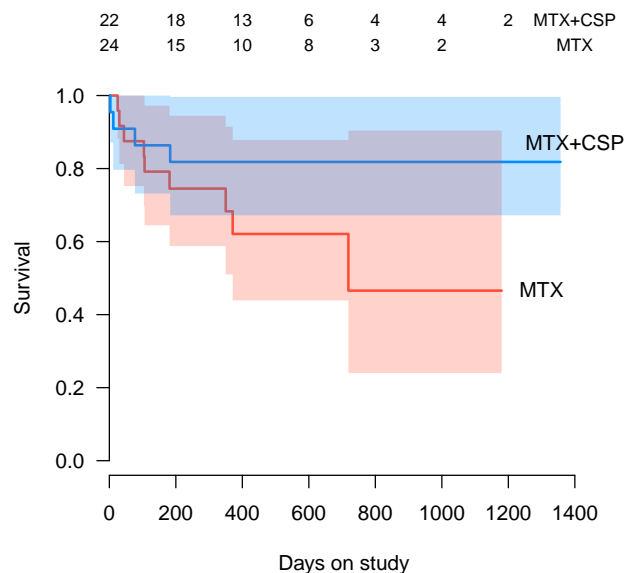


So as we can see, the median survival time for the MTX+CSP group never drops below 50%, so we can't estimate a median for it (nor do either of its confidence limits). Likewise, we can't estimate an upper confidence limit for the median survival in the MTX group, because its upper confidence limit (the dashed red line) never reaches 50%.

1.2 Fancier curves

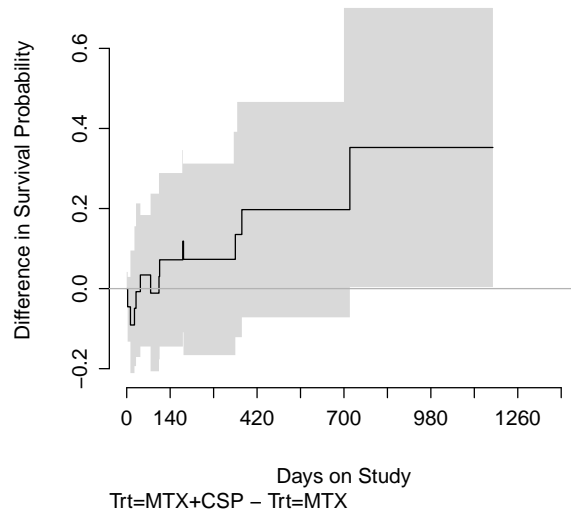
Personally, I greatly prefer shaded bands to separate dotted lines in illustrating confidence bands. I also like to have the option of putting the sample size (or number at risk) at the top of the plot, so I've written two functions that I personally use to add those things, called `ciband` and `nrisk`:

```
> source("http://myweb.uiowa.edu/pbreheny/5710/f16/labs/ciband.R")
> source("http://myweb.uiowa.edu/pbreheny/5710/f16/labs/nrisk.R")
> col <- c("#FF4E37FF", "#008DFFFF")
> plot(fit, xlab="Days on study", ylab="Survival", mark.time=FALSE,
+       col=col, lwd=2, las=1, bty="n")
> ciband(fit, col=c("#FF4E3740", "#008DFF40"))
> nrisk(fit)
> text(1400, 0.87, "MTX+CSP", xpd=TRUE)
> text(1300, 0.47, "MTX")
```



There are also several packages that can make fancier Kaplan-Meier plots like the one above. One in particular that is worth mentioning is `rms`, which can also make neat plots of the differences in survival. As we've said many times, just because confidence intervals overlap doesn't mean there isn't a significant difference between the groups:

```
> require(rms)
> fit.mort <- npsurv(S.mort ~ Trt)
> survdiffplot(fit.mort, xlab="Days on Study", order=2:1)
```



2 Log-rank tests

As we touched upon in class, the most common approach to testing for differences in survival is something called a log-rank test, which constructs contingency tables for each observed event time and then tests for differences by aggregating these tables.

```
> survdiff(S.mort ~ Trt)

Call:
survdiff(formula = S.mort ~ Trt)

          N Observed Expected (O-E)^2/E (O-E)^2/V
Trt=MTX   24         9   6.45    1.007    2.01
Trt=MTX+CSP 22         4   6.55    0.992    2.01

Chisq= 2 on 1 degrees of freedom, p= 0.156

> survdiff(S.gvhd ~ Trt)

Call:
survdiff(formula = S.gvhd ~ Trt)

          N Observed Expected (O-E)^2/E (O-E)^2/V
Trt=MTX   24        13   7.95    3.21    6.51
Trt=MTX+CSP 22         3   8.05    3.17    6.51

Chisq= 6.5 on 1 degrees of freedom, p= 0.0107
```

These tests indicate that in terms of mortality, we see a few more deaths than we would expect in the MTX group (9 vs. 6.5), but this could happen fairly easily by chance alone ($p = 0.16$). However, in terms of GVHD, we see quite a few more cases in the MTX group than we would expect (13 vs. 8), and this is not

so easily explained by chance ($p = 0.01$), suggesting that CSP is important for warding off GVHD. Whether it also improves survival or not is inconclusive, but it certainly doesn't seem to hurt.

This concludes our brief introduction to survival analysis; for more information, consider taking Survival Data Analysis (BIOS 7210) or Applied Survival Analysis (BIOS 6210).