Bayesian Survival Analysis – A Reality Check

Leonhard Held

University of Zurich[UZH]

O-Bayes19, University of Warwick, 30 June 2019

Variable selection in survival analysis

Central for the development of clinical prediction models:

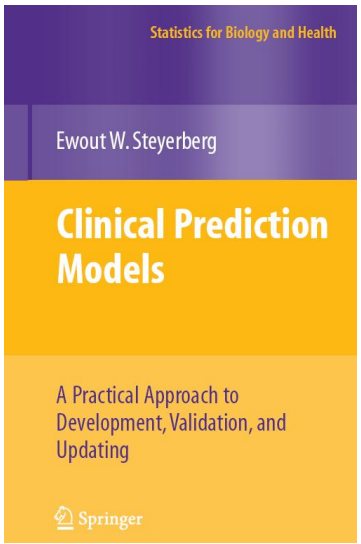

(2009)

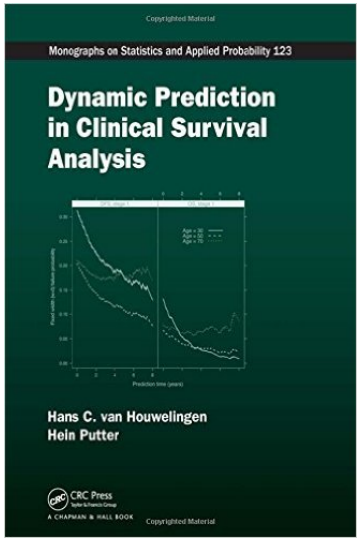

(2012)

More difficult than in standard linear model due to

- Right-censoring and left-truncation
- Time-dependent covariates or covariate effects
- Landmarking
- Competing risks

My own work in this area

Statistical Science
2015, Vol. 30, No. 2, 242–257
DOI: 10.1214/14-STS510

Approximate Bayesian Model Selection with the Deviance Statistic

Leonhard Held, Daniel Sabanés Bové and Isaac Gravestock

Statistics in Medicine

Research Article

Received 15 July 2015, Accepted 8 August 2016 Published online 31 August 2016 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.7089

Objective Bayesian model selection for Cox regression

Leonhard Held,[a*†] Isaac Gravestock[a] and Daniel Sabanés Bové[b]

Bayesian variable selection for censored survival data

Castellanos et al. (2019) consider the accelerated failure time model with log-normal errors and note:

> *Admittedly, the resulting methodology is computationally very demanding so, as it is, it's only suitable for small to moderate p.*

→ Let's try this with test-based Bayes factors based on the LR test statistic (Johnson, 2008; Hu and Johnson, 2009) with an implicit g-prior on the regression coefficients.

Different options to estimate g are described in Held et al. (2015, Stat Sci).

- See also Li and Clyde (2018, JASA) for additional justification.

Some background from the linear model

- Under the g-prior, the Bayes factor of a linear model $\mathcal{M}_j$ (with p additional parameters) vs. the null model $\mathcal{M}_0$ is

$$\mathrm{BF}_{j,0} = (g+1)^{(n-p-1)/2}\{1+g(1-R^2)\}^{-(n-1)/2}$$

  where R^2 is the coefficient of determination.
- Posterior mean is weighted average of MLE $\widehat{\beta}_{ML}$ and prior mean $\mathbf{0}$ with prior sample size $n/g \rightarrow$ shrinkage factor $t = g/(g+1)$
- The empirical Bayes estimate of g is

$$\hat{g}_{\mathrm{EB}} = \max\{F-1, 0\} \qquad \hat{t}_{\mathrm{EB}} = \max\{1-1/F, 0\}$$

  where F is the F-statistic of the global F test (Copas, 1983, 1997).

And for GLM's and other "regular" models

- Under the generalized g-prior, the test-based Bayes factor of a model $\mathcal{M}_j$ (with p additional parameters) vs. the null model $\mathcal{M}_0$ is

$$\mathrm{TBF}_{j,0} = \frac{\mathrm{p}(z \,|\, \mathcal{M}_j)}{\mathrm{p}(z \,|\, \mathcal{M}_0)} = (g+1)^{-p/2} \exp\left(\frac{g}{g+1}\frac{z}{2}\right)$$

  where z is the deviance.

→ n/g is still the prior sample size.

→ $t = g/(g+1)$ still serves as approximate shrinkage factor for $\widehat{\beta}_{ML}$.

- The empirical Bayes estimate of g is

$$\hat{g}_{\mathrm{EB}} = \max\{z/p - 1, 0\}$$

- The corresponding shrinkage factor is $\hat{t}_{\mathrm{EB}} = \max\{1 - p/z, 0\}$, as also proposed by Copas (1983) for optimal prediction in GLMs.

Step 1: Compute deviance for each model considered

Use function `survreg()` in `R`

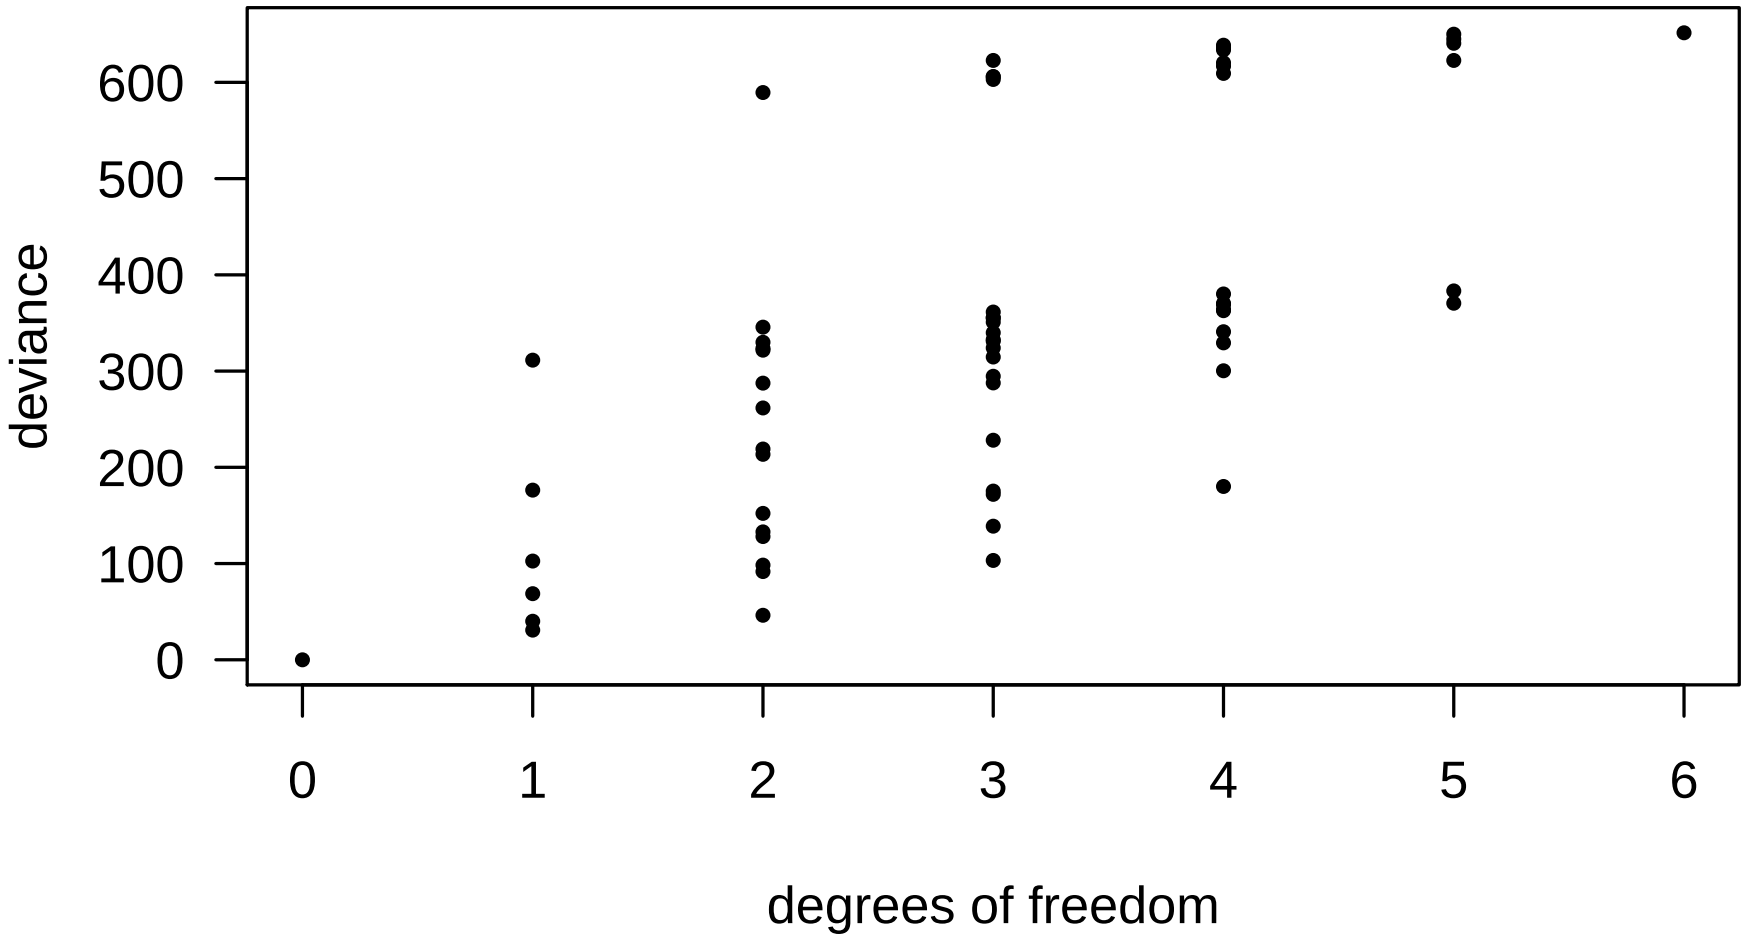

Step 2: Compute test-based Bayes factor

Based on the empirical Bayes estimate of g

```
## EB estimate of g
ghat <- pmax(ifelse(df>0, deviance/df - 1, 0), 0)
## Test-based Bayes factor to null model
TBF <- function(g, deviance, df){
    term <- ifelse(df>0, (g+1)^(-df/2)*exp(g/(g+1)*deviance/2), 1)
    return(term)
}
tbf1 <- TBF(ghat, deviance, df)
head(tbf1)

## [1] 1.000000e+00 1.151952e+21 8.781381e+36 6.904443e+59 1.406030e+66
## [6] 2.449711e+72
```

Step 3: Compute posterior probabilities

With multiplicity-corrected model prior

```
k <- 6
prior <- 1/(k+1)/choose(k, apply(I, 1, sum)-1)

post1 <- modelPosterior(prior, tbf1, order=TRUE)
head(post1)
```

```
##      nodes age metasta recurrence ER PGR    postProb
## [1,]     1   1       1          1  1   1 0.503531174
## [2,]     1   1       1          1  1   0 0.446580762
## [3,]     1   1       1          1  0   1 0.036001994
## [4,]     1   1       1          0  1   0 0.007279569
## [5,]     1   1       1          0  1   1 0.004043426
## [6,]     1   1       1          1  0   0 0.001894980
```

```
iP1 <- inclusionProbs(post1)
round(iP1, 2)
```

```
##      nodes        age    metasta recurrence         ER        PGR
##       1.00       1.00       1.00       0.99       0.96       0.54
```

Step 4: Compute posterior distribution

Apply shrinkage and compute model averages

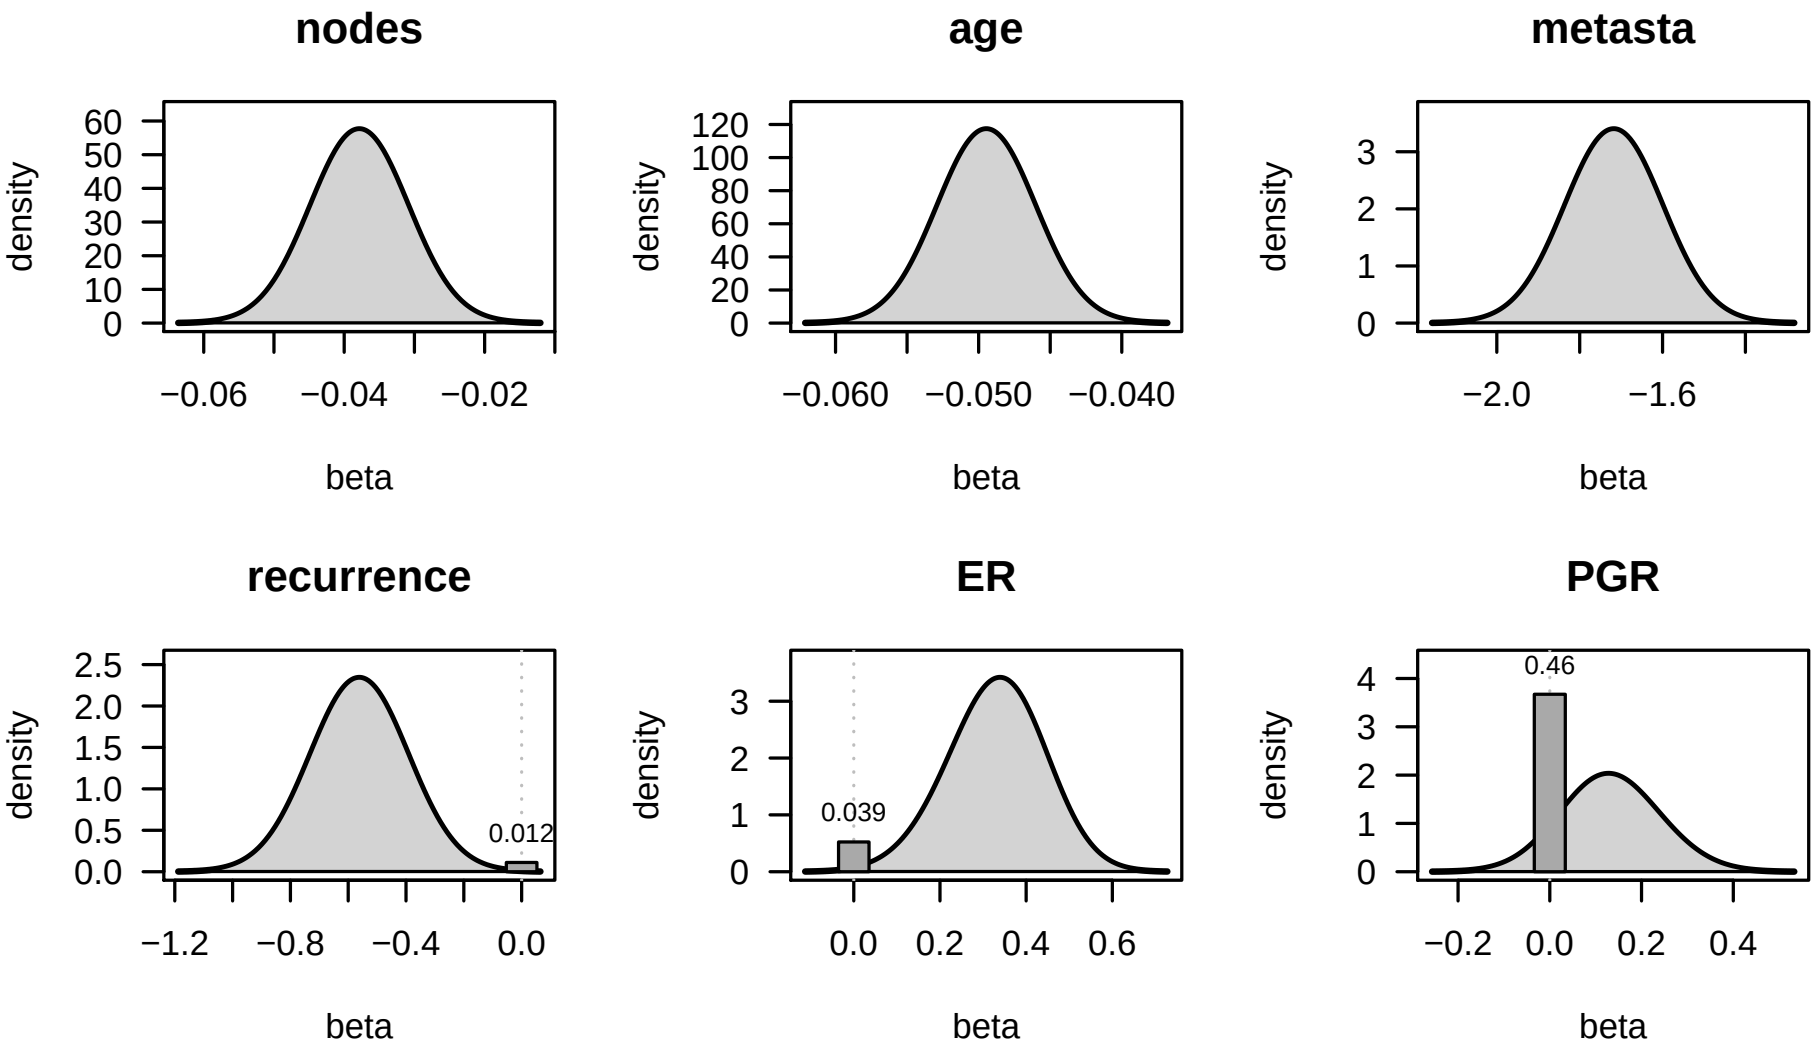

Practical implications

Q: Do we really need a "strict Bayesian perspective"?

- Prior on g can easily be incorporated in TBF approach.
- No problem at all to handle factors with TBFs.
- No problem to compute (model averaged) survival probabilities for specific patient covariates.
- Approach can be easily extended to larger p, dynamic variable selection, competing risks, etc

The role of the number of events n_u

- Standard error of the log hazard ratio is $\approx 2/\sqrt{n_u}$
- Sample size planning for a survival study is done in two steps:
  1. Determine the number of events n_u
  2. The required total number of individuals n is then calculated via

$$n = n_u / \Pr(\text{death})$$

  based on the assumed survival functions in the two groups and assumptions on the accrual and follow-up period.
- This has led Volinsky and Raftery (2000) to replace n with n_u in BIC obtained from a Cox model.
- Similiarly, the hyper-g/n prior is replaced by the hyper-g/n_u prior in Held et al. (2016, Stat Med).

Q: How sensitive are the results on the effective sample size with respect to the chosen parametric model?

Outlook: Dynamic variable selection for landmarking

Development of a clinical prediction model in the SMART cohort study

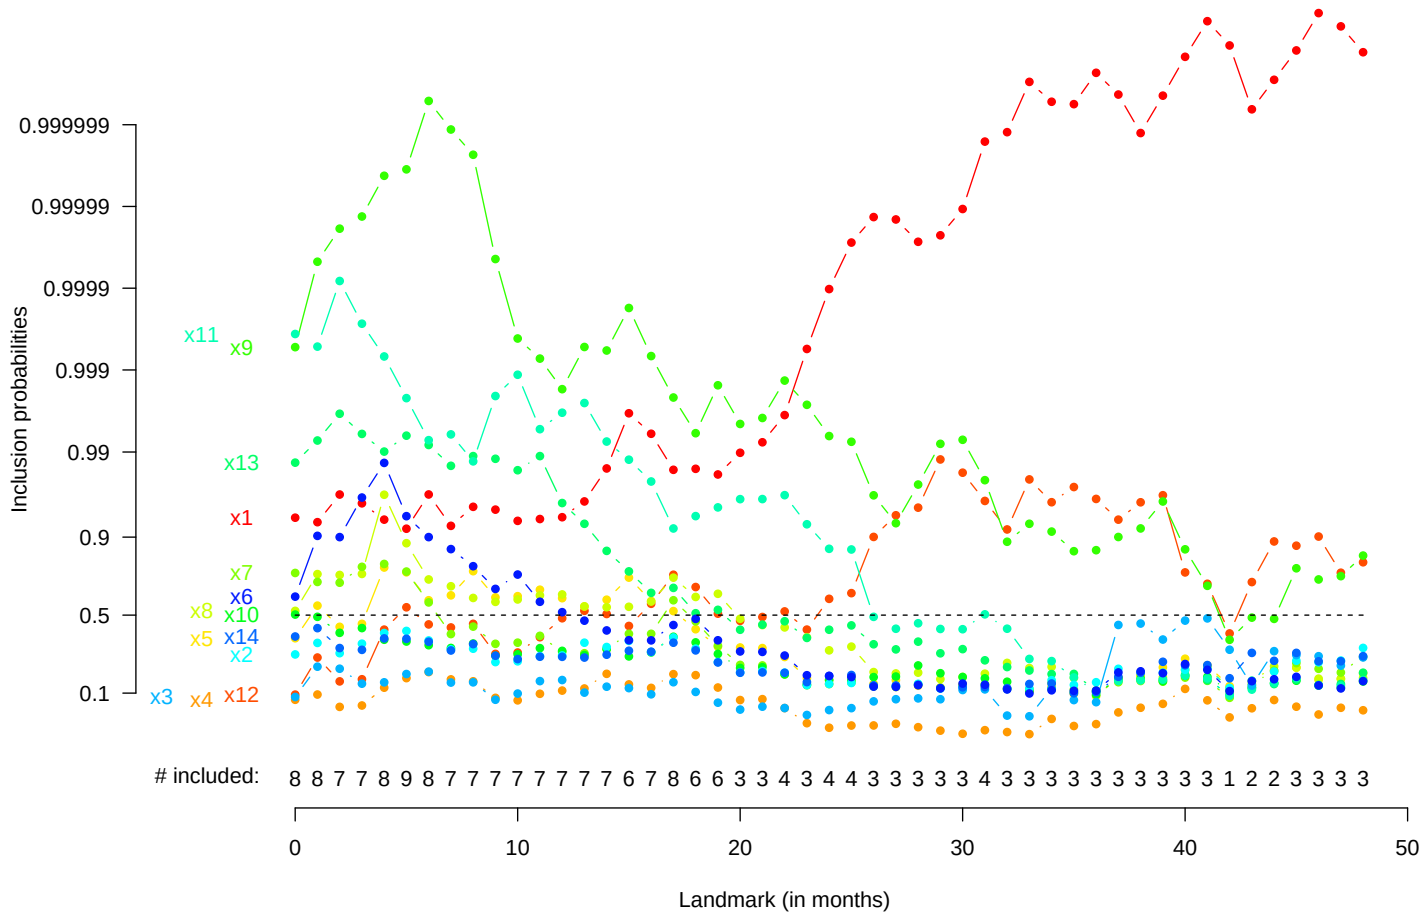

Parameter estimates

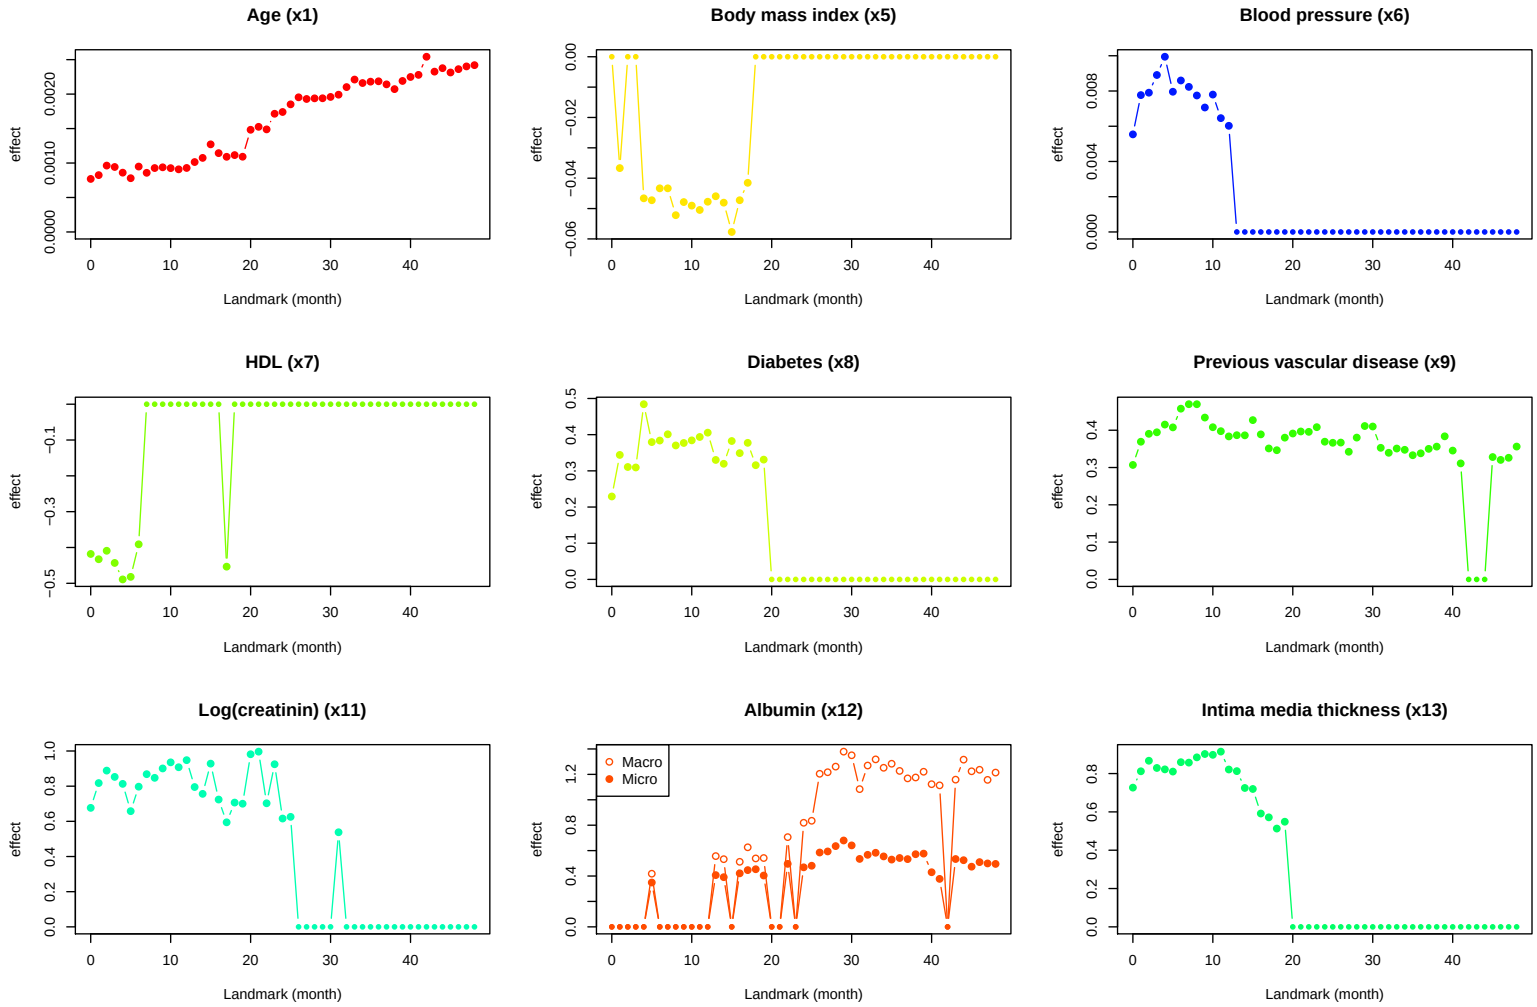

Outlook: Bayesian variable selection for competing risks

Received: 15 November 2017 | Revised: 10 August 2018 | Accepted: 14 August 2018

DOI: 10.1002/bimj.201700259

RESEARCH PAPER

Biometrical Journal

Dynamic clinical prediction models for discrete time-to-event data with competing risks—A case study on the OUTCOMEREA database

Rachel Heyard[1] | Jean-François Timsit[2] | Wafa Ibn Essaied[2] | Leonhard Held[1] |

on behalf of the COMBACTE-MAGNET consortium

OUTCOMEREA database

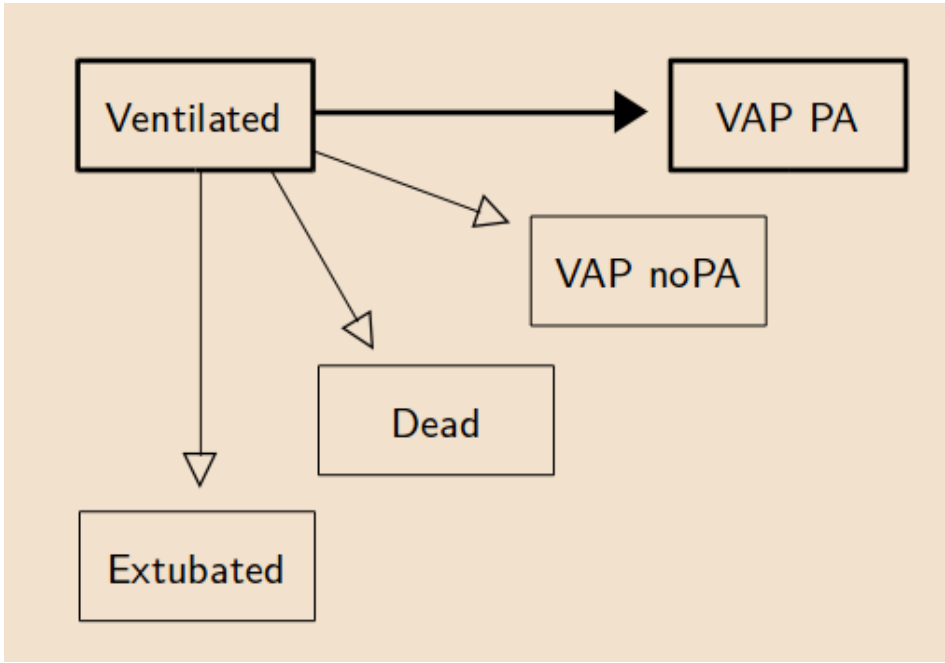

VAP: ventilator-associated pneumonia PA: Pseudomonas aeruginosa

OUTCOMEREA database with **daily** records on ICU admissions ($n = 7319$, $p = 26$)

Cause-specific variable selection

Standardized coefficients

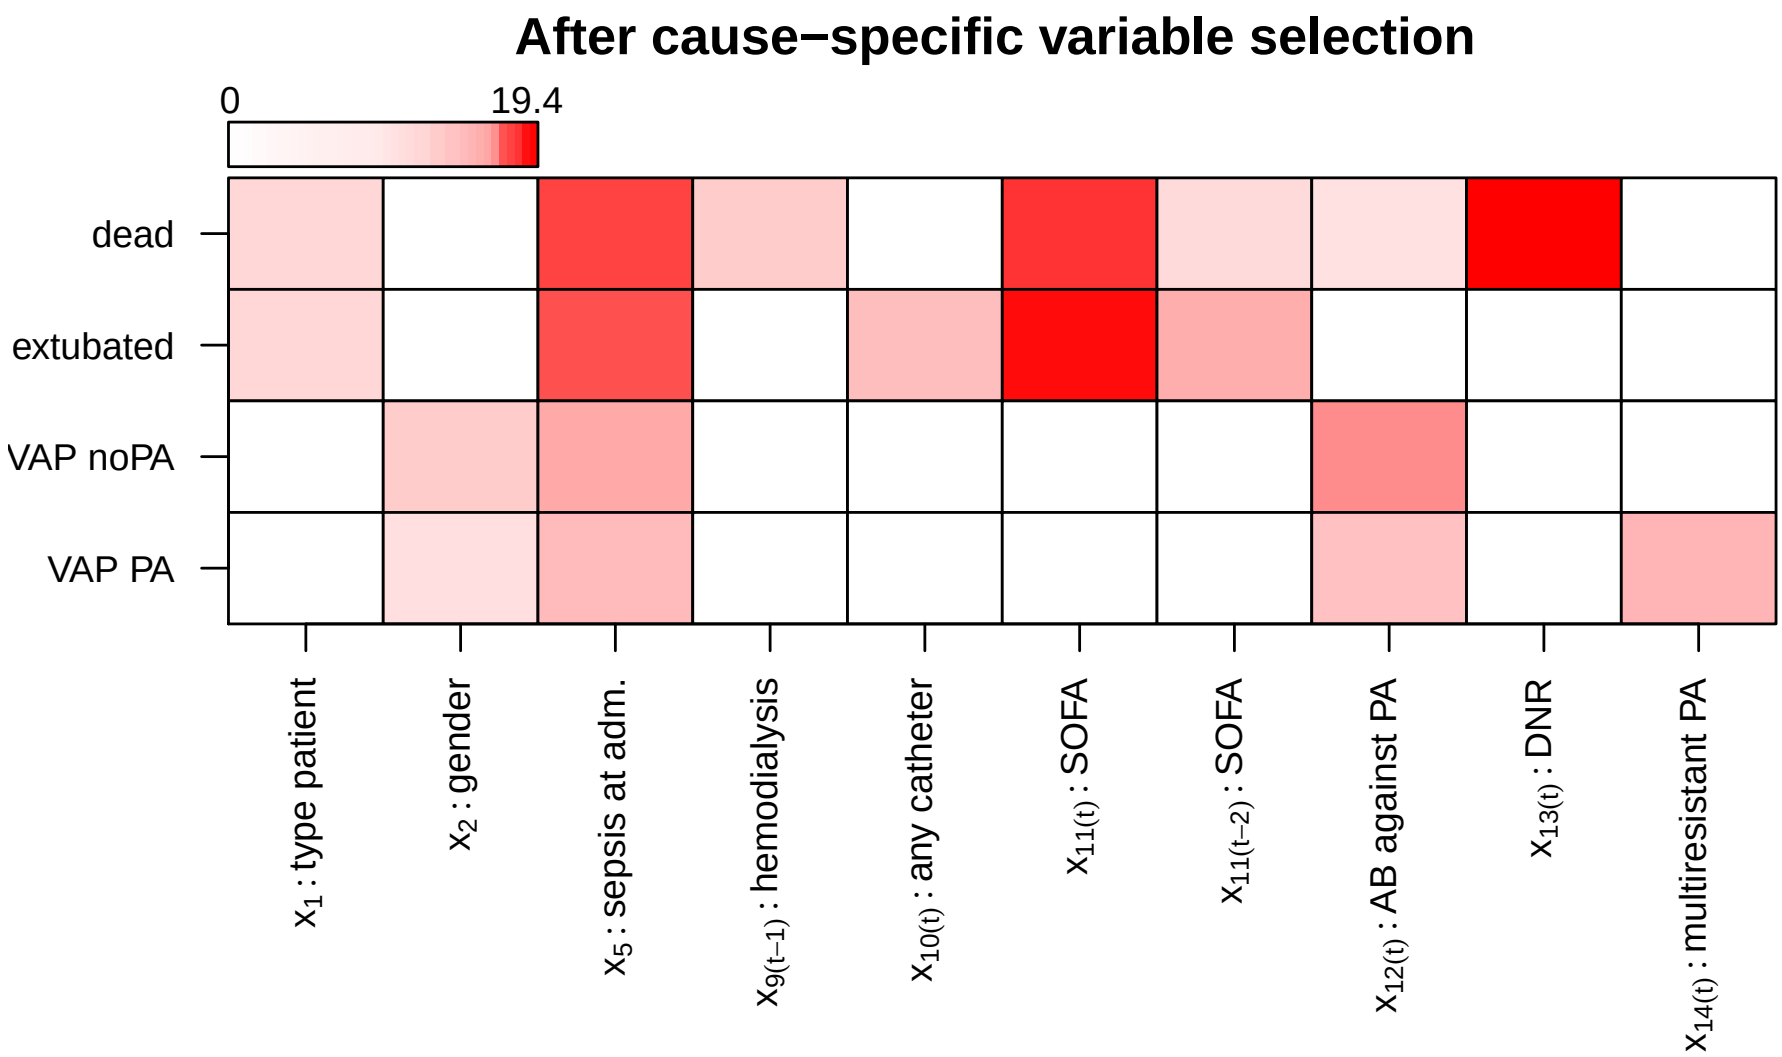

Uncertainty quantification for the survival curve

Method by Link (1984)

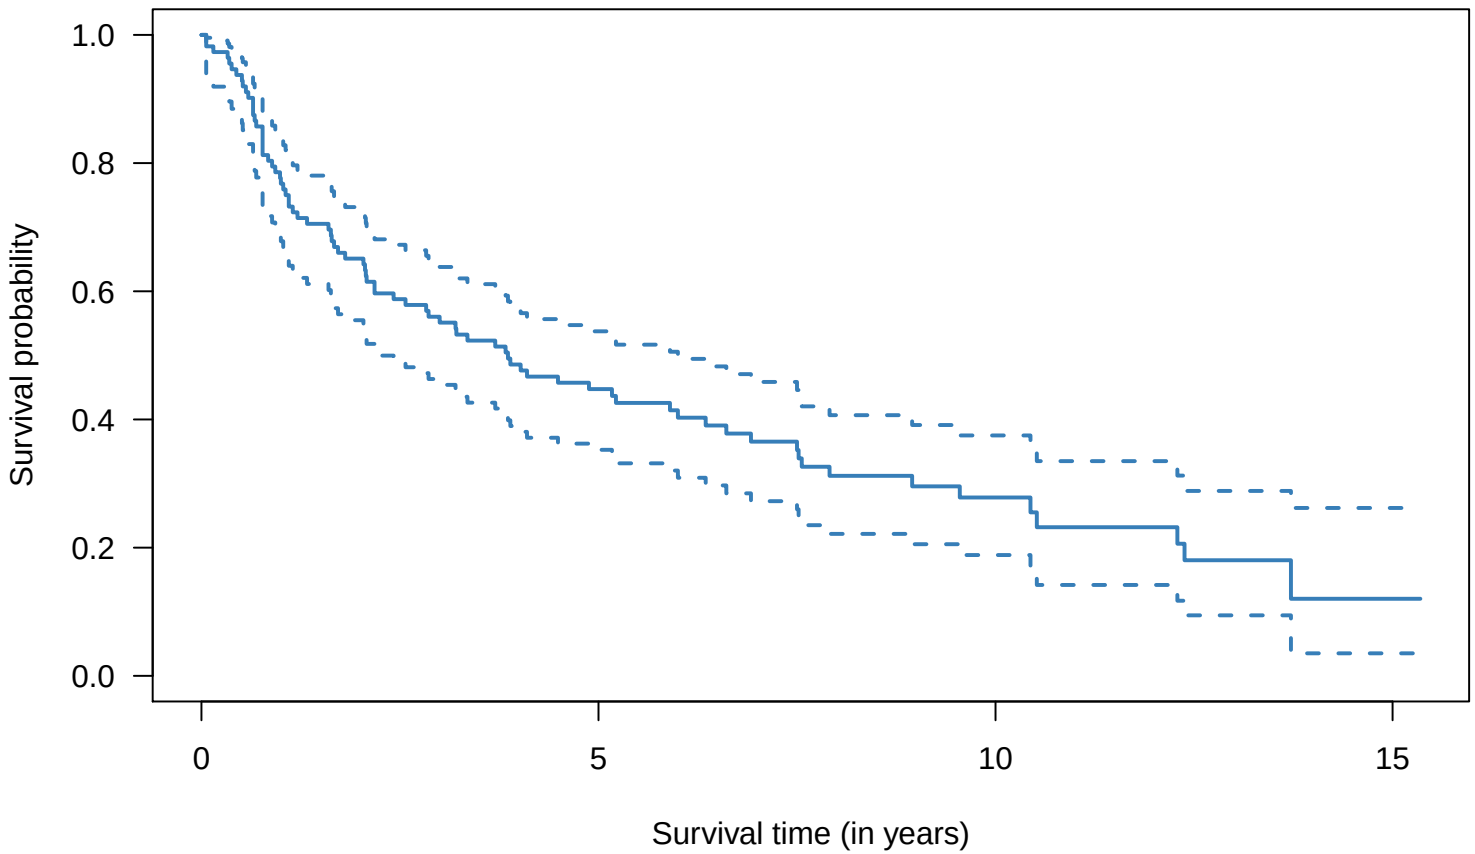

The median survival time

Method by Brookmeyer and Crowley (1982)

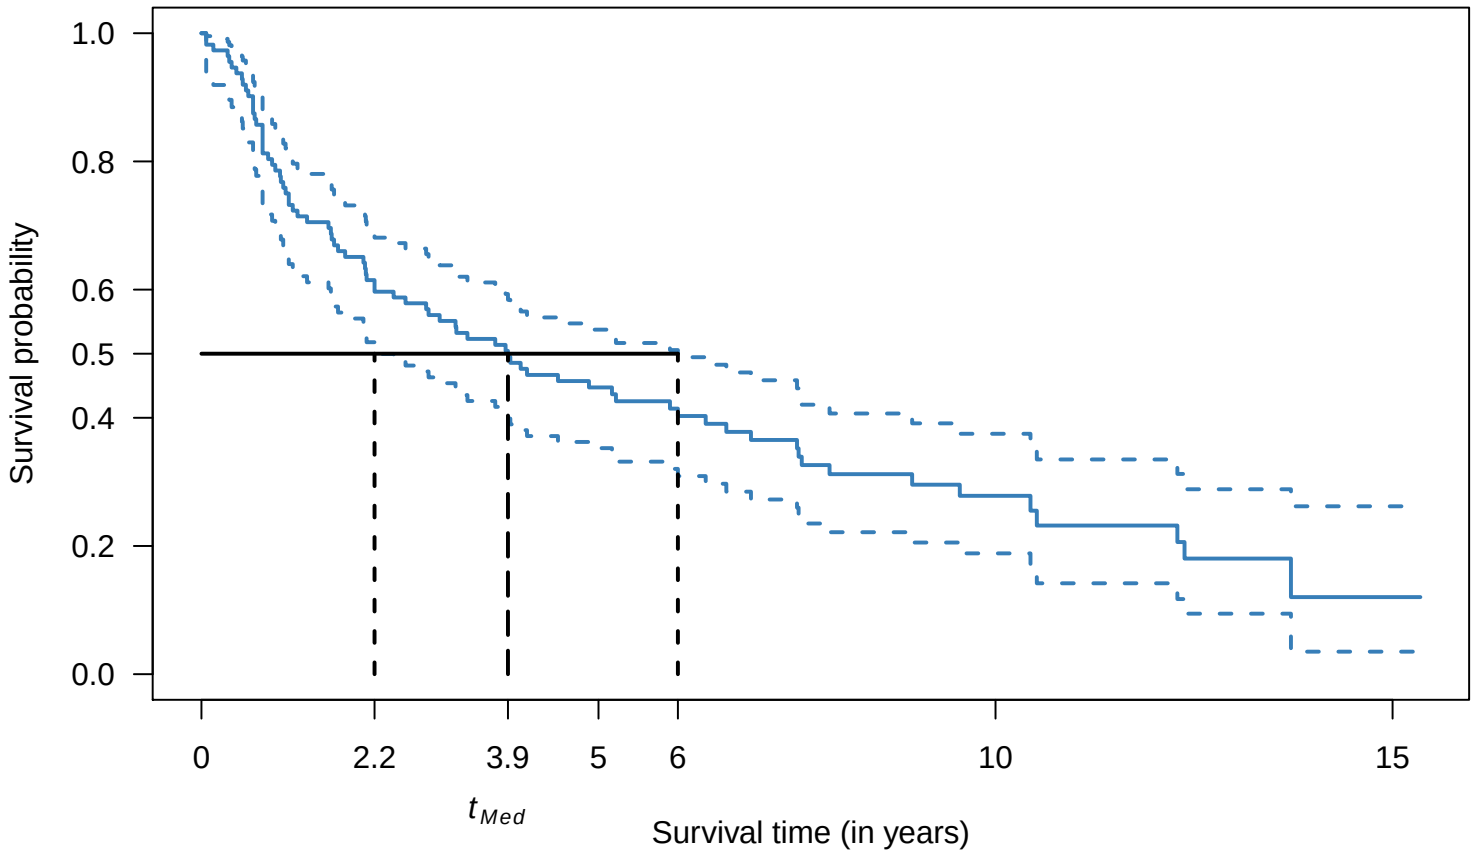

Q: Can you provide credible intervals for the median survival time?

The role of the number of intervals K

- Castillo and van der Pas (2019) assume piecewise-constant hazard functions.
- In the application the number of intervals is $K = 10$.
- Similiarly, the life table method assumes piecewise-constant hazards.
- Kaplan-Meier is a limiting case where each interval contains only one event (in the absence of ties).

Q: Can you do the same to free yourself from arbitrarily fixing K?

Literature I

Brookmeyer, R. and Crowley, J. (1982). A confidence interval for the median survival time. *Biometrics*, 38(1):29.

Castellanos, M. E., Garcia-Donato, G., and Cabras, S. (2019). Variable selection priors for survival models with censored data. Technical report.

Copas, J. B. (1983). Regression, prediction and shrinkage. *Journal of the Royal Statistical Society. Series B (Methodological)*, 45(3):311–354.

Copas, J. B. (1997). Using regression models for prediction: shrinkage and regression to the mean. *Statistical Methods in Medical Research*, 6(2):167–183.

Held, L., Gravestock, I., and Sabanés Bové, D. (2016). Objective bayesian model selection for cox regression. *Statistics in Medicine*, 35(29):5376–5390.

Held, L., Sabanés Bové, D., and Gravestock, I. (2015). Approximate Bayesian model selection with the deviance statistic. *Statistical Science*. to appear.

Hu, J. and Johnson, V. E. (2009). Bayesian model selection using test statistics. *Journal of the Royal Statistical Society. Series B (Methodological)*, 71(1):143–158.

Johnson, V. E. (2008). Properties of Bayes factors based on test statistics. *Scandinavian Journal of Statistics*, 35(2):354–368.

Li, Y. and Clyde, M. A. (2018). Mixtures of g-priors in generalized linear models. *Journal of the American Statistical Association*, 113(524):1828–1845.

Literature II

Link, C. L. (1984). Confidence intervals for the survival function using cox's proportional-hazard model with covariates. *Biometrics*, 40(3):601.

Volinsky, C. T. and Raftery, A. E. (2000). Bayesian information criterion for censored survival models. *Biometrics*, 56:256–262.