

Survival Prediction Using Gene Expression Data

Zhu Wang
UT Health San Antonio

Abstract

This document describes applications of R package **bujar** for predicting survival in diffuse large B cell lymphoma treated with chemotherapy plus Rituximab using gene expression data.

Keywords: survival, Buckley-James regression, prediction, boosting, variable selection.

1. Introduction

Researchers have been interested in predicting survival in diffuse large B cell lymphoma (DLBCL) treated with chemotherapy. We re-evaluate clinical and microarray data in [Lenz et al. \(2008\)](#). Data from two treatment plans were collected: CHOP and R-CHOP. CHOP is a combination chemotherapy with cyclophosphamide, doxorubicin, vincristine and prednisone. The current gold standard includes rituxima immunotherapy in addition to the chemotherapy (R-CHOP). It was interesting to identify genes that predict survival among patients who received CHOP also retain their prognostic power among patients who received R-CHOP. Data from 181 CHOP patients (training data) are used to build predictive models and data from 233 R-CHOP patients (test data) are used to validate the models. Due to the nature of high-dimensional data with 54675 probe sets or covariates, we first conduct a pre-selection procedure on the training data. This procedure filters out genes with lower variations if a sample variance for a gene was smaller than the 10th percentile for that gene. Details on how the data were read and pre-screened may be found in `dlbcl_raw.R` in the vignettes directory. Afterwards we have 3833 remaining probe sets which were stored in the R package **bujar** ([Wang 2015](#)). Test data with the same remaining genes is used for validation.

One of the challenges is that we have right censored data, as often occurred in survival data analysis. Although proportional-hazards regression is a popular option, we focus on Buckley-James (BJ) regression for high-dimensional data proposed by [Wang and Wang \(2010\)](#). Essentially, BJ regression iteratively imputes the censored survival outcomes using the conditional expectation obtained from the current estimates, and the imputed survival outcomes are re-fitted with regression techniques, which in turn lead to different BJ predictive models depend on what regression model is chosen. We evaluate BJ boosting algorithms and BJ penalized regression implemented in **bujar**. To avoid overfitting, predictive modeling typically requires tuning parameter(s) selection. Tuning parameter is selected by data-driven cross-validation, unless otherwise specified. However, if twin boosting is conducted, the boosting iteration number is fixed in the second round of boosting. The analysis results can be different from [Wang and Wang \(2010\)](#) for two reasons. First, computer codes have been changed for BJ boost-

ing with componentwise least squares or smoothing splines as base learner. In the current implementation since version 0.1-10, every BJ iteration uses the optimal boosting iteration obtained from the last BJ iteration if tuning parameter is selected. The change results in more sparse variable selection and also consistent with BJ boosting with trees as base learner. Second, BJ boosting (not twin boosting) with trees involves an internal random mechanism for tree model building. In the current implementation since version 0.1-10, a new parameter `rng` in function `bujar` can be used for reproducible results.

```
R> library("bujar")
R> data("chop")
R> data("rchop")
R> ###censoring rate in the CHOP data
R> sum(chop$status==0)/nrow(chop)

[1] 0.4199

R> rchop <- subset(rchop, select=colnames(rchop)%in% colnames(chop))
R> chop$survtime <- chop$survtime + 1 ### add 1 for log-transformation
```

2. BJ Boosting with Linear Regression

2.1. BJ boosting with componentwise least squares (BJ-LS)

```
R> set.seed(123)
R> res.lin <- bujar(y=log(chop[,1]), cens=chop[,2], x=chop[,-(1:2)], tuning=TRUE,
  cv=TRUE, n.cores=1, mstop=1000)
R> ###number of genes selected with BJ-LS
R> sum(res.lin$xselect==1)
R> coef.bj <- coef(res.lin)
R> ###estimated non-zero coefficients (only list 10)
R> coef.bj[abs(coef.bj)>0][1:10]

R> library("survival")
R> cutyear <- 3

R> pred.bj <- predict(res.lin, newx=rchop[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff
```

```
R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))
```

Figure 1: Kaplan-Meier survival curves for BJ-LS regression.

2.2. BJ twin boosting with componentwise least squares

```
R> res.lin2 <- bujar(y=log(chop[,1]), cens=chop[,2], x=chop[,-(1:2)], tuning=TRUE,
  cv=FALSE, mstop=1000, twin=TRUE, mstop2=100)
R> ### number of genes selected with BJ-LS
R> sum(res.lin2$xselect==1)
R> coef.bj <- coef(res.lin2)
R> coef.bj[abs(coef.bj)>0]
R> pred.bj <- predict(res.lin2, newx=rchop[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))
```

Figure 2: Kaplan-Meier survival curves for twin BJ-LS regression.

To reduce computing burden for other modeling methods, we adopted a supervised gene screening to select the top 100 probe sets based on univariate BJ regression.

```
R> library("rms")
R> res <- rep(NA,ncol(chop))
R> for(i in 3:ncol(chop)){
  bjres <- try(bj(Surv(survtime, status) ~ chop[,i],data=chop, link="log"))
  ###if BJ convergence fails, still included for further analysis
  if(inherits(bjres, "try-error")) res[i] <- 1e-5
  else res[i] <- anova(bjres)[1,3] #p-value
}
R> nsel <- 100
R> ### select top nsel=100 genes with most significant p-values
R> chop2 <- chop[, c(1, 2, sort.list(res,decreasing=FALSE)[1:nsel])]
R> rchop2 <- rchop[, c(1, 2, sort.list(res,decreasing=FALSE)[1:nsel])]
R> colnames(chop2)[- (1:2)] <- colnames(rchop2)[- (1:2)] <-
  paste("x",colnames(chop2)[- (1:2)],sep="")
R> detach(package:rms)
```

3. BJ LASSO

Within each BJ iteration, the LASSO is used to fit the imputed survival outcomes. The penalty tuning parameter is fixed at the 20th value in 100 penalty sequence values determined within each BJ iteration.

```
R> res.lasso <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="enet2", tuning=FALSE, whichlambda=20)
R> ### how many genes selected by BJ-LASSO
R> sum(res.lasso$xselect==1)
R> ###estimated non-zero coefficients (only list 10)
R> coef.bj <- coef(res.lasso)
R> coef.bj[abs(coef.bj)>0][1:10]
R> pred.bj <- predict(res.lasso, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))
```

Figure 3: Kaplan-Meier survival curves for BJ-LASSO regression

4. BJ SCAD

BJ-SCAD is an extension of BJ-LASSO (Wang and Wang 2010; Fan and Li 2001). Within each BJ iteration, the SCAD is used with the imputed survival outcomes. The penalty tuning parameter is fixed at the 20th value in 100 penalty sequence values determined within each BJ iteration.

```
R> res.scad <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="snet", tuning=FALSE, whichlambda=20)
R> ### how many genes selected by BJ-SCAD
R> sum(res.scad$xselect==1)
R> ###estimated non-zero coefficients (only list 10)
R> coef.bj <- coef(res.scad)
R> coef.bj[abs(coef.bj)>0][1:10]
R> pred.bj <- predict(res.scad, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))
```

Figure 4: Kaplan-Meier survival curves for BJ-SCAD regression.

5. BJ Boosting with Smoothing Splines

5.1. BJ boosting with componentwise smoothing splines (BJ-SM)

```

R> set.seed(123)
R> res.ss <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="pspline", tuning=FALSE, cv=FALSE, mstop=100)
R> ### how many genes selected by BJ smoothing splines, only list 10
R> sum(res.ss$xselect==1)
R> colnames(res.ss$x)[res.ss$xselect==1][1:10]
R> pred.bj <- predict(res.ss, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))

```

Figure 5: Kaplan-Meier survival curves for BJ-SM regression.

5.2. BJ twin boosting with componentwise smoothing splines

```

R> set.seed(123)
R> res.ss2 <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="pspline", tuning=TRUE, cv=TRUE, n.cores=1, mstop=100, twin=TRUE, mstop2=200)
R> ### how many genes selected by BJ twin smoothing splines, only list 10
R> sum(res.ss2$xselect==1)
R> colnames(res.ss2$x)[res.ss2$xselect==1][1:10]
R> pred.bj <- predict(res.ss2, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))

```

Figure 6: Kaplan-Meier survival curves for twin BJ-SM regression.

6. BJ Boosting with Regression Trees

6.1. BJ boosting with regression stumps (BJ-Tree)

```
R> res.tree1 <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="tree",tuning=TRUE, cv=TRUE, mstop=1000, n.cores=1, rng=123)
R> ###Number of genes selected with tree, only list 10
R> sum(res.tree1$xselect==1)
R> colnames(res.tree1$x)[res.tree1$xselect==1][1:10]
R> pred.bj <- predict(res.tree1, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))
```

Figure 7: Kaplan-Meier survival curves for BJ-Tree regression.

6.2. BJ twin boosting with regression stumps

```

R> res.tree2 <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="tree", tuning=TRUE, cv=TRUE, mstop=1000, twin=TRUE, mstop2=100,
  n.cores=1, rng=123)
R> ###Number of genes selected with tree, only list 10
R> sum(res.tree2$xselect==1)
R> colnames(res.tree2$x)[res.tree2$xselect==1][1:10]
R> pred.bj <- predict(res.tree2, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))

```

Figure 8: Kaplan-Meier survival curves for twin BJ-Tree regression.

6.3. BJ boosting with regression trees of degree 4

```

R> res.tree4 <- bujar(y=log(chop2[,1]), cens=chop2[,2], x=chop2[,-(1:2)],
  learner="tree",degree=4, tuning=TRUE, cv=TRUE, mstop=100, rel.inf=TRUE,
  n.cores=1,rng=123)
R> ###Number of genes selected with tree, only list 10
R> sum(res.tree4$xselect==1)
R> colnames(res.tree4$x)[res.tree4$xselect==1][1:10]
R> pred.bj <- predict(res.tree4, newx=rchop2[,-(1:2)])
R> pred.bj <- exp(pred.bj) - 1
R> group <- cut(pred.bj, breaks=c(-1, cutyear, 100), labels=c("high", "low"))
R> dat.km <- data.frame(survtime=rchop$survtime, status = rchop$status, group=group)
R> fit.diff <- survdiff(Surv(survtime, status) ~ group, data=dat.km)
R> fit.diff

R> fit.surv <- survfit(Surv(survtime, status) ~ group, data=dat.km )
R> plot(fit.surv, xlab="Year past therapy",ylab="Survival probability",
  lty = 1:2, col=c("red","blue"), mark.time=TRUE)
R> legend(1, .1, c("High risk", "Low risk"), lty = 1:2, col=c("red","blue"))

```

Figure 9: Kaplan-Meier survival curves for BJ-Tree (degree 4) regression.

Partial dependence plots can be utilized to show the impact of one or more covariates on the response after taking account the average effects of all other covariates in the model. We can generate plots for selected genes. See Figure 5 in [Wang and Wang \(2010\)](#).

```
R> gene <- c("x1558999_x_at", "x212713_at", "x224043_s_at", "x229839_at",
  "x237515_at", "x237797_at", "x242758_x_at", "x244346_at")
R> library("gridExtra")
R> for(i in 1:length(gene))
  eval(parse(text=(paste("a", i, " <- plot(res.tree4$res.fit,
    i.var=which(colnames(res.tree4$x) == gene[i])), sep=""))))
R> grid.arrange(a1, a2, a3, a4, a5, a6, a7, a8, ncol=4)
```

Figure 10: Partial plots of selected genes based on BJ-Tree (degree=4) regression.

The two-way interaction partial plots display gene-gene interactions similar to Figure 6 in [Wang and Wang \(2010\)](#).

```
R> for(i in 1:6)
  eval(parse(text=(paste("b", i, " <- plot(res.tree4$res.fit,
                                i.var=unlist(res.tree4$interactions$rank.list[i,c(1, 3)]))",
                                sep=""))))
R> grid.arrange(b1, b2, b3, b4, b5, b6, ncol=2)
```

Figure 11: Gene-gene interactions based on BJ-Tree (degree=4) regression.

7. SessionInfo

```
R> sessionInfo();
```

```
R version 3.6.1 (2019-07-05)
```

```
Platform: x86_64-pc-linux-gnu (64-bit)
```

```
Running under: Ubuntu 18.04.3 LTS
```

```
Matrix products: default
```

```
BLAS: /usr/lib/x86_64-linux-gnu/blas/libblas.so.3.7.1
```

```
LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.7.1
```

```
locale:
```

```
[1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C               LC_TIME=en_US.UTF-8
[4] LC_COLLATE=C              LC_MONETARY=en_US.UTF-8    LC_MESSAGES=en_US.UTF-8
[7] LC_PAPER=en_US.UTF-8      LC_NAME=C                  LC_ADDRESS=C
[10] LC_TELEPHONE=C            LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
```

```
attached base packages:
```

```
[1] stats      graphics  grDevices  utils      datasets  methods    base
```

```
other attached packages:
```

```
[1] survival_2.44-1.1 bujar_0.2-7
```

```
loaded via a namespace (and not attached):
```

```
[1] splines_3.6.1      foreach_1.4.4      R.utils_2.8.0      elasticnet_1.1.1
[5] Formula_1.2-3      assertthat_0.2.1   stats4_3.6.1       latticeExtra_0.6-28
[9] stabs_0.6-3        numDeriv_2016.8-1 pillar_1.3.1       backports_1.1.3
[13] lattice_0.20-38    quantreg_5.38      glue_1.3.1         quadprog_1.5-5
[17] mboost_2.9-1       digest_0.6.18      RColorBrewer_1.1-2 checkmate_1.9.1
[21] sandwich_2.5-0     colorspace_1.4-0   gbm_2.1.5          rms_5.1-3
[25] htmltools_0.3.6    Matrix_1.2-17      R.oo_1.22.0        plyr_1.8.4
[29] pkgconfig_2.0.2    SparseM_1.77       earth_4.7.0        purrr_0.3.0
[33] mvtnorm_1.0-8      scales_1.0.0       TeachingDemos_2.10 MatrixModels_0.4-1
[37] tibble_2.1.3       htmlTable_1.13.1   mda_0.4-10         ggplot2_3.1.0
[41] TH.data_1.0-10     nnet_7.3-12        lazyeval_0.2.1     magrittr_1.5
[45] crayon_1.3.4       polyspline_1.1.13  R.methodsS3_1.7.1  nlme_3.1-141
[49] R.cache_0.13.0     doParallel_1.0.14 MASS_7.3-51.4      foreign_0.8-72
[53] class_7.3-15       R.rsp_0.43.1       tools_3.6.1        data.table_1.12.0
[57] multcomp_1.4-8     stringr_1.4.0      munsell_0.5.0      glmnet_2.0-16
[61] cluster_2.1.0      plotrix_3.7-4      compiler_3.6.1     inum_1.0-0
[65] rlang_0.4.0        plotmo_3.5.2       grid_3.6.1         iterators_1.0.10
[69] rstudioapi_0.9.0   mpath_0.3-17       htmlwidgets_1.3    base64enc_0.1-3
[73] partykit_1.2-3     gtable_0.2.0       codetools_0.2-16   lars_1.2
[77] R6_2.4.0           zoo_1.8-4          bst_0.3-18         gridExtra_2.3
[81] nnls_1.4           knitr_1.21         dplyr_0.8.3        Hmisc_4.2-0
```

[85]	libcoin_1.0-3	modeltools_0.2-22	stringi_1.3.1	pscl_1.5.2
[89]	parallel_3.6.1	Rcpp_1.0.1	rpart_4.1-15	acepack_1.4.1
[93]	tidyselect_0.2.5	xfun_0.4		

References

- Fan J, Li R (2001). “Variable selection via nonconcave penalized likelihood and its oracle properties.” *Journal of the American Statistical Association*, **96**(456), 1348–1360.
- Lenz G, Wright G, Dave SS, Xiao W, Powell J, Zhao H, Xu W, Tan B, Goldschmidt N, Iqbal J, Vose J, Bast M, Fu K, Weisenburger DD, Greiner TC, Armitage JO, Kyle A, May L, Gascoyne RD, Connors JM, Troen G, Holte H, Kvaloy S, Dierickx D, Verhoef G, Delabie J, Smeland EB, Jares P, Martinez A, Lopez-Guillermo A, Montserrat E, Campo E, Braziel RM, Miller TP, Rimsza LM, Cook JR, Pohlman B, Sweetenham J, Tubbs RR, Fisher RI, Hartmann E, Rosenwald A, Ott G, Muller-Hermelink H, Wrench D, Lister TA, Jaffe ES, Wilson WH, Chan WC, Staudt LM (2008). “Stromal gene signatures in large-B-cell lymphomas.” *New England Journal of Medicine*, **359**(22), 2313–2323.
- Wang Z (2015). *bujar: Buckley-James regression for survival data with high-dimensional covariates*. R package version 0.2-5, URL <http://CRAN.R-project.org/package=bujar>.
- Wang Z, Wang CY (2010). “Buckley-James boosting for survival analysis with high-dimensional biomarker data.” *Statistical Applications in Genetics and Molecular Biology*, **9**(1), Article 24.

Affiliation:

Zhu Wang
Department of Epidemiology & Biostatistics
UT Health San Antonio
E-mail: wangz1@uthscsa.edu