Using glmnetr

Walter K. Kremers, Mayo Clinic, Rochester MN

18 February 2023

The Package

For some datasets, for example when there are collinearities in the design matrix x, *glmnet* may have very long run times when fitting the relaxed lasso model, making it difficult to get solutions either from *cv.glmnet* or even *glmnet*. In this package, *glmnetr*, we provide a workaround and solve for the unpenalized relaxed model where gamma=0 for the linear, logsitc and Cox regression model structures using the *stats* glm() function and the R *survival* pacakge coxph() function. If you are not fitting relaxed lasso models, or if you are able to reasonably quickly get convergence using *glmnet*, then this package may not be of much benefit to you. Note, while this package may allow one to fit relaxed lasso models that have difficulties converging using *glmnet*, this package does not afford the full function and versatility of *glmnet*.

In addition to fitting the relaxed lasso model this package also includes the function cv.glmnetr() to perform a cross validation (CV) to identify hyperparameters for a lasso fit, much like the cv.glmnet() function of the *glmnet* package. Additionally, the package includes the function nested.glmnetr() to perform a nested CV to assess the fit of a cross validation informed lasso model fit. If though you are fitting not a relaxed lasso model but an elastic-net model, then the R-packages *nestedcv* (https://cran.r-project.org/package=nestedcv), 'glmnetSE' (https://cran.r-project.org/package=glmnetSE) or others may provide greater functionality when performing a nested CV.

As with the *glmnet* package, this package passes most relevant information to the output object which can be evaluated using plot, summary and predict functions. Use of the *glmnetr* package has many similarites to the *glmnet* package and it is recommended that the user of *glmnetr* first become familiar with the *glmnet* package (https://cran.r-project.org/package=glmnet), with the "An Introduction to glmnet" and "The Relaxed Lasso" being especially helpful in this regard.

Data requirements

The basic data elements for input to the *glmnetr* analysis programs are similar to those of *glmnet* and include 1) a matrix of predictors and 2) an outcome variable or variables in vector form. For the estimation of the "fully" relaxed models (where gamma=0) the package is set up to fit the "gaussian" and "binomial" models using the *stats* glm() function and Cox survival models using the the coxph() function of the *survival* pacakge. When fitting the Cox model the outcome model variable is interpreted as the "time" variable in the Cox model, and one must also specify 3) a variable for event, agian in vector form, and optionally 4) a variable for start time, also in vector form. Row i of the predictor matrix and element i of the outcome vector(s) are to include the data for the same sampling unit.

An example dataset

To demonstrate usage of *glmnetr* we first generate a data set for analysis, run an analysis and evaluate using the plot(), summary() and predict() functions.

The code

```
# Simulate data for use in an example relaxed lasso fit of survival data
# first, optionally, assign a seed for random number generation to get applicable results
set.seed(116291949)
simdata=glmnetr.simdata(nrows=1000, ncols=100, beta=NULL)
```

generates simulated data for analysis. We extract data in the format required for input to the *glmnetr* programs.

```
# Extract simulated survival data
xs = simdata$xs         # matrix of predictors
y_ = simdata$yt         # vector of survival times
event = simdata$event   # indicator of event vs. censoring
```

Inspecting the predictor matrix we see

```
# Check the sample size and number of predictors
print(dim(xs))
```

```
## [1] 1000  100
```

```
# Check the rank of the design matrix, i.e. the degrees of freedom in the predictors
rankMatrix(xs)[[1]]
```

```
## [1] 94
```

```
# Inspect the first few rows and some select columns
print(xs[1:10,c(1:12,18:20)])
```

```
##       X1 X2 X3 X4 X5 X6 X7 X8 X9 X10 X11 X12        X18        X19         X20
##  [1,]  1  1  0  0  0  0  0  0  0   1   0   1  0.1513225 -0.4034383  0.35250844
##  [2,]  1  0  0  0  1  0  0  1  0   0   0   0 -1.1610480  0.5533030  0.14578868
##  [3,]  1  0  0  1  0  0  1  0  0   0   0   0 -0.3292269  0.3086399 -0.48443836
##  [4,]  1  1  0  0  0  0  0  0  0   1   0   0  2.0635214 -0.5500741 -0.02173104
##  [5,]  1  0  0  0  1  0  0  1  0   0   0   0  0.3905722 -0.6836452 -0.37643201
##  [6,]  1  0  1  0  0  0  0  0  1   0   0   0 -0.2397597  1.6909447  0.49599945
##  [7,]  1  0  1  0  0  0  0  1  0   0   0   0 -0.5592424  0.2314638 -0.53198341
##  [8,]  1  0  0  1  0  0  0  0  0   0   1   0 -1.0050514  0.5319574  0.54287646
##  [9,]  1  0  0  1  0  0  0  0  0   0   1   0  1.2548034  0.8213164  0.17067691
## [10,]  1  0  0  0  1  0  0  0  1   0   0   0 -0.3079151 -0.6105910 -0.88711869
```

Cross validation (CV) informed relaxed lasso model fit

To fit a relaxed lasso model and get reasonable hyperparameters for lambda and gamma, and summarize the cross-validated "tuned" model fit, we can use the function cv.glmnet() and summary() functions.

```
# Fit a relaxed lasso model informed by cross validation
cv.cox.fit = suppressWarnings( cv.glmnetr(xs, NULL, y_, event, family="cox", track=0) )
```

Note, in the derivation of the relaxed lasso model fits, individual coefficients may be unstable even when the model may be stable which elicits warning messages. Thus we "wrapped" the call to cv.glmnetr() within the suppressWarnings() function to suppress excessive warning messages in this vignette. The first term in the call to cv.glmnetr(), xs, is the design matrix for predictors. The second input term, here NULL, is for the start time in case (start, stop) time data setup is used in a Cox survival model. The third term is the outcome variable for the linear regression or logistic regression model and the time of event or censoring in case of the Cox model, and finally the forth term is the event indicator variable for the Cox model taking the value 1 in case of an event or 0 in case of censoring at time y_. The forth term would be NULL for either linear or logistic regression. Currently the options for family are "guassian" for linear regression , "binomial" for logistic regression (both using the *stats* glm() function) and "cox" for the Cox proportional hazards regression model using the coxph() function of the R *survival* package. If one sets track=1 the program will update progres in the R console, else for track=0 it will not. | Before numerically summarizing the model fit, or inspecting the coefficient estimates, we plot the average deviances using the plot function.

```
# Plot cross validation average deviances for a relaxed lasso model
plot(cv.cox.fit)
```

```
##  min CV average deviance (max log likelihood) for
##    relaxed at log(lambda)  = -3.937, gamma.min = 0.75, df = 36
##    fully relaxed at log(lambda)   = -3.472, df = 18
##    fully penalized at log(lambda) = -4.216, df = 48
```

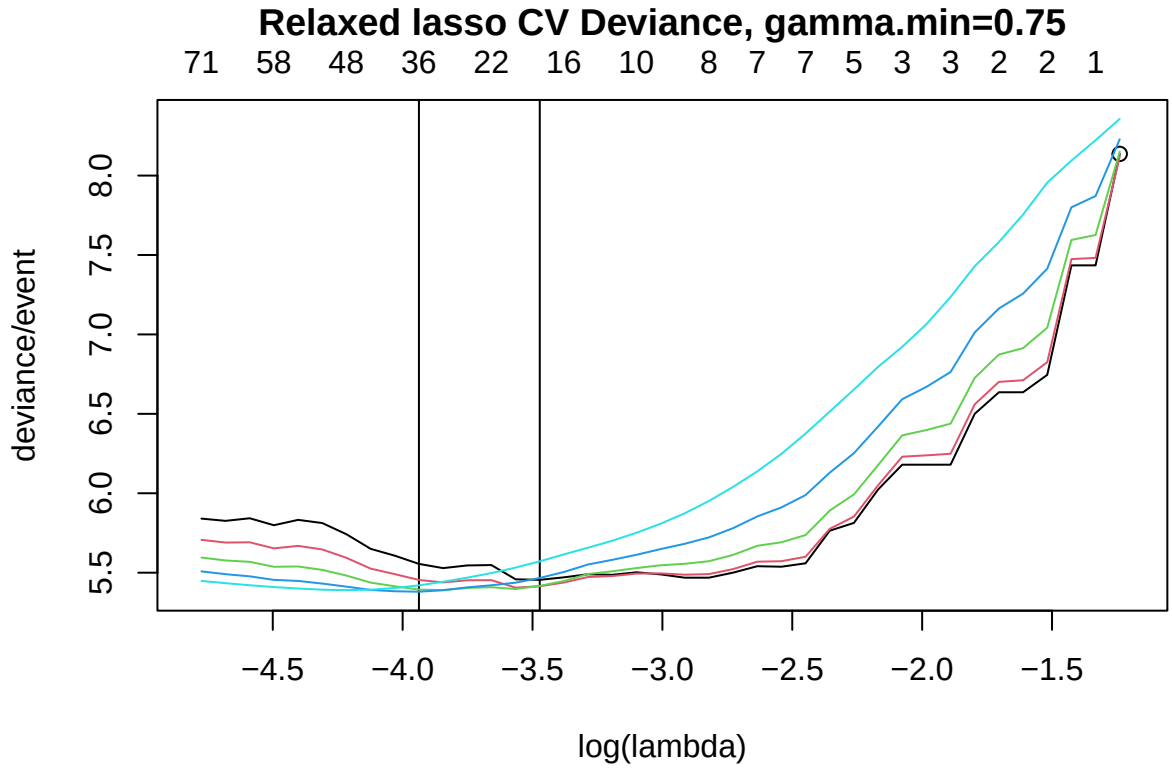

In that to maximize the log-likelihoods is to minimize deviance we inspect these curves for a minimum. The minimizing lambda is indicated by the left most vertical line, here about log(lambda) = -3.94. The minimizing gamma is 0.75 and described in the title. Whereas there is no legend here for gamma, when non-zero coefficients start to enter the model as the penalty is reduced, here shown to the right, deviances

will tend to be smaller for gamma = 0, greater for gamma = 1 and in between for other gammas values. From this figure we also see that at lambda=0.75 the deviance is hardly distinguishable for gamma ranging from 0.5 to 1. More relevant we see that the fully relaxed lasso (gamma=0) and indicated by the right most vertical line, achieves a "nearly" minimal deviance at about -3.47.

```
# Plot coefficients informed by a cross validation
plot(cv.cox.fit, coefs=TRUE)
```

```
##  min CV average deviance (max log likelihood)
##   at log(lambda.min) = -3.937, gamma.min = 0.75, df = 36
```

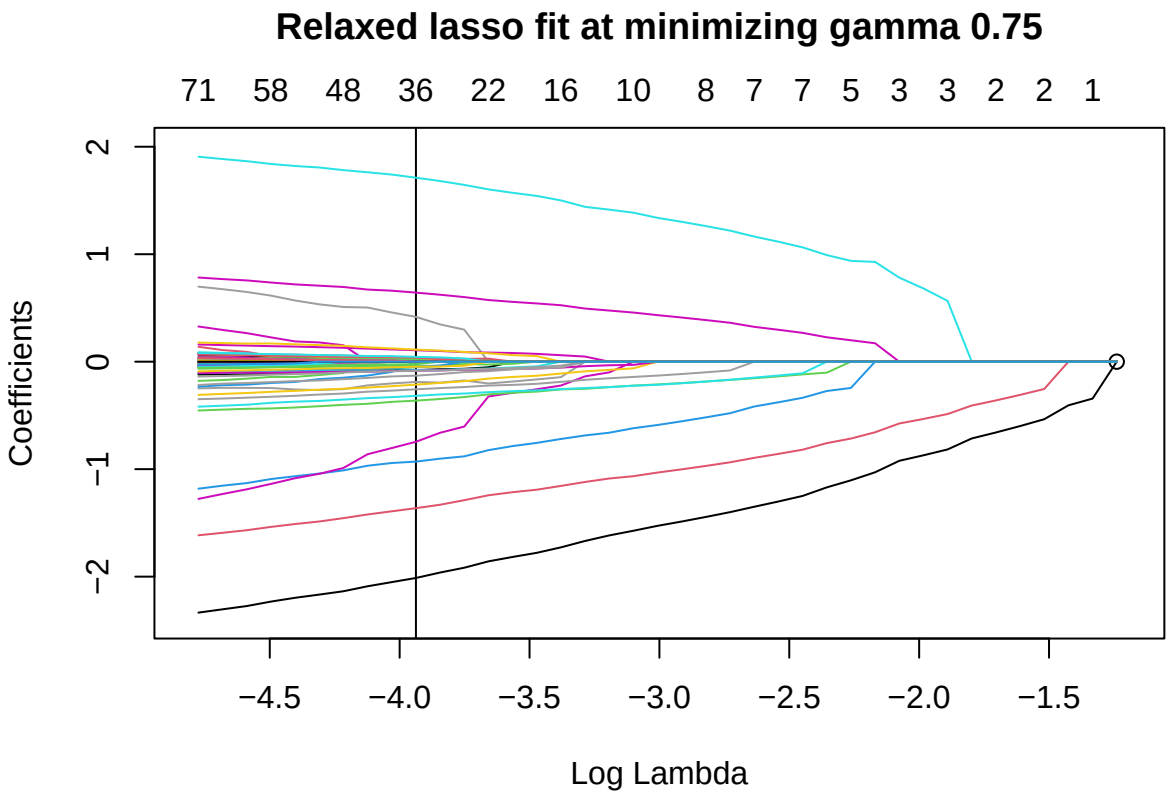

In this plot of coefficients we use the same orientation for lambda as in the plot for deviances with larger values of the lambda penalty to the right and corresponding to fewer non-zero coefficients. The displayed coefficients are for the minimizing gamma=0.75 as noted in the tile, and the minimizing lambda indicated by the vertical line. Now, since the fully relaxed lasso model had a deviance almost that of the relaxed lasso model we also plot the coefficients using the option gam=0.

```
# Plot fully relaxed coefficients informed by a cross validation
plot(cv.cox.fit, coefs=TRUE, gam=0)
```

```
##  Fully relaxed min CV average deviance (max log likelihood)
##   at log(lambda.min) = -3.472, df = 18
```

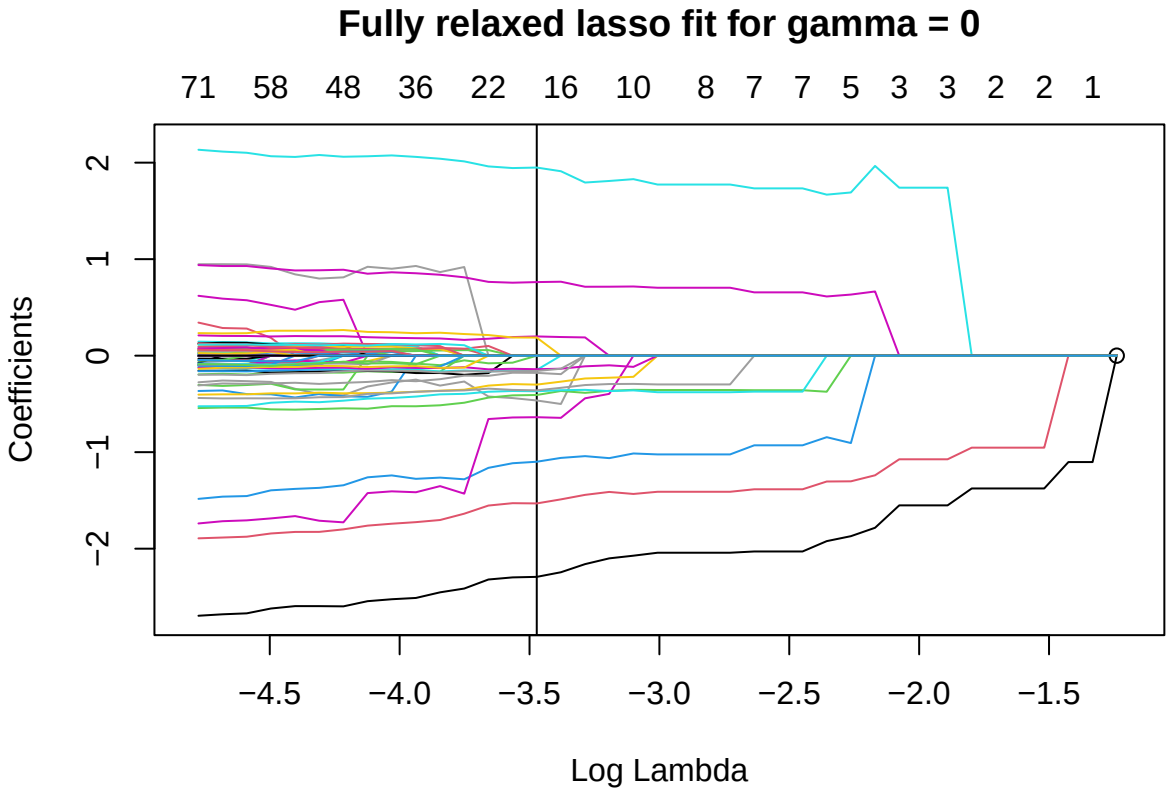

In addition to simply showing how the coefficients change as the lambda penalty is decreased, this plot shows how the coefficients change for the un-penalized (fully relaxed) model with gamma=0 as lambda decreases. In particular we see the coefficients become slightly larger in magnitude as the lambda penalty decreases and also as additional terms come into the model. This is not unexpected as omitted terms from the Cox model tend to bias coefficients toward 0 more than increase the standard error. We also see, as too indicated in the deviance plot, the number of model non-zero coefficients, 18, to be substantially less than the 36 from the relaxed lasso fit and the 48 from the fully penalized lasso fit.
| The summary function describes the relaxed lasso fit informed by CV.

```
# Summarize relaxed lasso model fit informed by cross validation
summary(cv.cox.fit)
```

```
##
##   The relaxed minimum is obtained for lambda = 0.01950023 , index = 30 and gamma = 0.75
##   with df (number of non-zero terms) = 36, average deviance = 5.380322 and beta =
##            X4            X5            X8           X14           X15
## -9.298178e-01  1.711721e+00 -1.885456e-01 -7.454261e-01  1.190965e-14
##           X16           X17           X18           X19           X20
##  4.167433e-01 -1.302087e-15 -1.362967e+00 -3.625131e-01  2.922682e-02
##           X21           X22           X23           X24           X25
## -3.182677e-01  6.426824e-01 -2.134638e-01 -2.569618e-01 -2.012441e+00
##           X29           X30           X34           X35           X38
##  2.389336e-02  2.816325e-02  2.602897e-02  2.119959e-02  1.064479e-01
##           X39           X43           X49           X50           X58
##  2.293174e-02 -5.218288e-02 -7.819052e-02  3.211853e-02  3.392804e-02
##           X59           X60           X61           X62           X67
## -3.948692e-02 -4.923017e-02 -7.808802e-02 -8.469542e-02  1.799604e-02
```

```
##           X71           X72           X77           X87           X96
## -4.703098e-02 -1.302645e-01  4.472362e-02  1.123160e-01 -8.482900e-02
##           X99
## -2.308179e-02
##
##   The fully relaxed (gamma=0) minimum is obtained for lambda = 0.03104988 and index = 25
##   with df (number of non-zero terms) = 17, average deviance = 5.455056 and beta =
##         X4         X5         X8        X14        X18        X19        X21
## -1.0991234  1.9492327 -0.4643855 -0.6377441 -1.5308945 -0.4075762 -0.3687472
##        X22        X23        X24        X25        X38        X61        X62
##  0.7625238 -0.2999622 -0.3593211 -2.2920690  0.1993155 -0.1477441 -0.1403819
##        X72        X87        X96
## -0.1784984  0.1874497 -0.1588458
##
##   The UNrelaxed (gamma=1) minimum is obtained for lambda = 0.01475121 and index = 33
##   with df (number of non-zero terms) = 48, average deviance = 5.390245
##
##
##   Order coefficients entered into the lasso model (1st to last):
##  [1] "X25" "X18" "X5"  "X22" "X4"  "X19" "X21" "X24" "X23" "X62" "X14" "X38"
## [13] "X8"  "X72" "X96" "X61" "X87" "X59" "X35" "X49" "X50" "X16" "X34" "X71"
## [25] "X77" "X30" "X39" "X43" "X58" "X60" "X20" "X29" "X67" "X99" "X12" "X28"
## [37] "X84" "X98" "X51" "X55" "X88" "X6"  "X11" "X33" "X78" "X40" "X66" "X68"
```

In the summary output we first see the relaxed lasso model fit based upon the (lambda, gamma) pair which minimizes the cross validated average deviance. Next is the model fit based upon the lambda that minimizes the cross validated average deviance along the path where gamma=0, that is among the fully relaxed lasso models. After that is information on the fully penalized lasso fit, but without the actual coefficient estimates. These estimates can be printed using the option *printg1=TRUE*, but are suppressed by default for space. Finally, the order that coefficients enter the lasso model as the penalty is decreased is provided, which gives some indication of relative model importance of the coefficients. Because, though, the differences in successive lambda values used in the numerical algorithms may allow multiple new terms to enter into the model between successive numerical steps, the ordering in this list may not be strict. If the user would want they could read lambda from output$lambda, set up a new lambda with finer steps and rerun the model. Our experience though is that this does not generally lead to a meaningfully different model and so is not done by default or as option. | One can as well use the predict function to get the coefficients for the lasso model, or the xs_new*beta for a new design matrix xs_new. In contrast to the summary function which simply displays coefficients, the predict function provides an outpout object in vector form (or a list with two vectors) and so can more easily be used for further calculations. By default the summary function will use the (lambda, gamma) pair that minimizes the average CV deviances. One can also specify lam=NULL and gam=1 to use the fully penalized lasso best fit, that use the solution that minimizes the CV deviance with respect to lambda while holding gamma=1, or gam=0 to use the fully relaxed lasso best fit, that is minimizes while holding gamma=0. One can also numerically specify both lam for lambda and gam for gamma. Within the package lambda and gamma usually denote vectors for the search algorithm and so other names are used uere.

```
# Get coefficients
beta = predict(cv.cox.fit)
```

```
##
##  (lambda, gamma) pair minimizing CV average deviance is used
```

```
# Print out the non-zero coefficients
beta$beta
```

```
##            X4            X5            X8           X14           X15
## -9.298178e-01  1.711721e+00 -1.885456e-01 -7.454261e-01  1.190965e-14
##           X16           X17           X18           X19           X20
##  4.167433e-01 -1.302087e-15 -1.362967e+00 -3.625131e-01  2.922682e-02
##           X21           X22           X23           X24           X25
## -3.182677e-01  6.426824e-01 -2.134638e-01 -2.569618e-01 -2.012441e+00
##           X29           X30           X34           X35           X38
##  2.389336e-02  2.816325e-02  2.602897e-02  2.119959e-02  1.064479e-01
##           X39           X43           X49           X50           X58
##  2.293174e-02 -5.218288e-02 -7.819052e-02  3.211853e-02  3.392804e-02
##           X59           X60           X61           X62           X67
## -3.948692e-02 -4.923017e-02 -7.808802e-02 -8.469542e-02  1.799604e-02
##           X71           X72           X77           X87           X96
## -4.703098e-02 -1.302645e-01  4.472362e-02  1.123160e-01 -8.482900e-02
##           X99
## -2.308179e-02
```

```
# Print out all coefficients
beta$beta_
```

```
##            X1            X2            X3            X4            X5
##  0.000000e+00  0.000000e+00  0.000000e+00 -9.298178e-01  1.711721e+00
##            X6            X7            X8            X9           X10
##  0.000000e+00  0.000000e+00 -1.885456e-01  0.000000e+00  0.000000e+00
##           X11           X12           X13           X14           X15
##  0.000000e+00  0.000000e+00  0.000000e+00 -7.454261e-01  1.190965e-14
##           X16           X17           X18           X19           X20
##  4.167433e-01 -1.302087e-15 -1.362967e+00 -3.625131e-01  2.922682e-02
##           X21           X22           X23           X24           X25
## -3.182677e-01  6.426824e-01 -2.134638e-01 -2.569618e-01 -2.012441e+00
##           X26           X27           X28           X29           X30
##  0.000000e+00  0.000000e+00  0.000000e+00  2.389336e-02  2.816325e-02
##           X31           X32           X33           X34           X35
##  0.000000e+00  0.000000e+00  0.000000e+00  2.602897e-02  2.119959e-02
##           X36           X37           X38           X39           X40
##  0.000000e+00  0.000000e+00  1.064479e-01  2.293174e-02  0.000000e+00
##           X41           X42           X43           X44           X45
##  0.000000e+00  0.000000e+00 -5.218288e-02  0.000000e+00  0.000000e+00
##           X46           X47           X48           X49           X50
##  0.000000e+00  0.000000e+00  0.000000e+00 -7.819052e-02  3.211853e-02
##           X51           X52           X53           X54           X55
##  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00
##           X56           X57           X58           X59           X60
##  0.000000e+00  0.000000e+00  3.392804e-02 -3.948692e-02 -4.923017e-02
##           X61           X62           X63           X64           X65
## -7.808802e-02 -8.469542e-02  0.000000e+00  0.000000e+00  0.000000e+00
##           X66           X67           X68           X69           X70
##  0.000000e+00  1.799604e-02  0.000000e+00  0.000000e+00  0.000000e+00
##           X71           X72           X73           X74           X75
## -4.703098e-02 -1.302645e-01  0.000000e+00  0.000000e+00  0.000000e+00
```

```
##           X76           X77           X78           X79           X80
##  0.000000e+00  4.472362e-02  0.000000e+00  0.000000e+00  0.000000e+00
##           X81           X82           X83           X84           X85
##  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00
##           X86           X87           X88           X89           X90
##  0.000000e+00  1.123160e-01  0.000000e+00  0.000000e+00  0.000000e+00
##           X91           X92           X93           X94           X95
##  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00  0.000000e+00
##           X96           X97           X98           X99          X100
## -8.482900e-02  0.000000e+00  0.000000e+00 -2.308179e-02  0.000000e+00
```

```
# Get the predicteds (linear predictors) for the original data set
predicteds = predict(cv.cox.fit, xs)
```

```
##
##  (lambda, gamma) pair minimizing CV average deviance is used
```

```
# Print out the first few predicteds
predicteds[1:20]
```

```
##  [1]  0.1009869  4.3123200 -4.4673094 -1.7814957  0.9077405 -1.0938730
##  [7]  4.6147550 -0.1989544 -6.0377112 -1.6112803 -0.8198837  2.1165700
## [13] -1.1446918 -2.1868484  1.7288336 -0.1977422 -1.0923555 -4.4471458
## [19]  2.2903899 -2.8503090
```

Model fit without cross validation

We can as well fit a relaxed lasso model without doing a CV. For this case one can still plot the coefficients but when the minimizing lambda and gamma are not informed by CV one is to specify which gamma should be used for the plots. By default gamma=1, i.e. for the fully penalized lasso model, is used for the plots. One can plot the coefficient estimates for different gamma values, but these will usually be more meaningful when informed by the CV "tuned" hyperparameters values for lambda and gamma. One can also use the predict() function, again to output either coefficients or predicteds, i.e. xs_new*beta for a new design matrix xs_new. Such predicteds are often, for example in coxph(), included in the analysis output object under the name linear.predictors.

```
# Fit a model without cross validation
cox.fit = suppressWarnings( glmnetr(xs, NULL, y_, event, family="cox") )
# Plot coefficients of the fully relaxed lasso model
plot(cox.fit, gam=0)
```

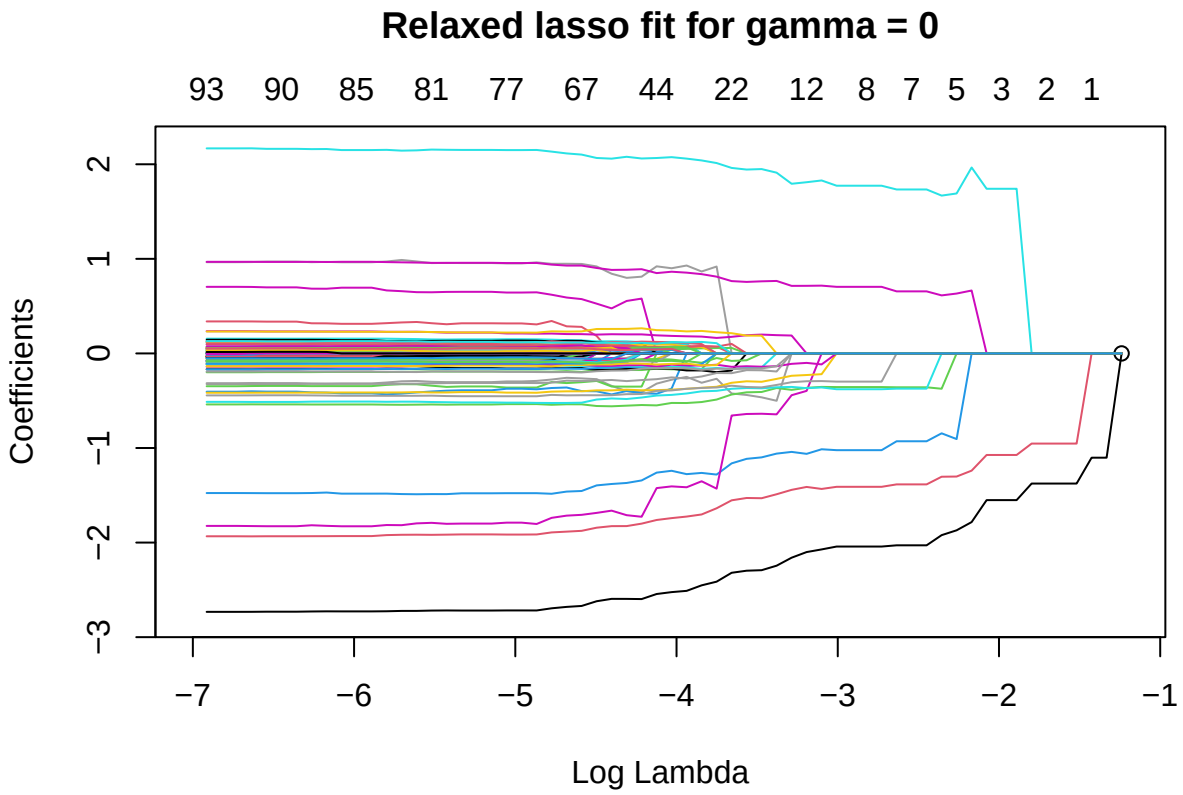


```
# Plot coefficients of the fully penalized lasso model
plot(cox.fit, gam=1)
```

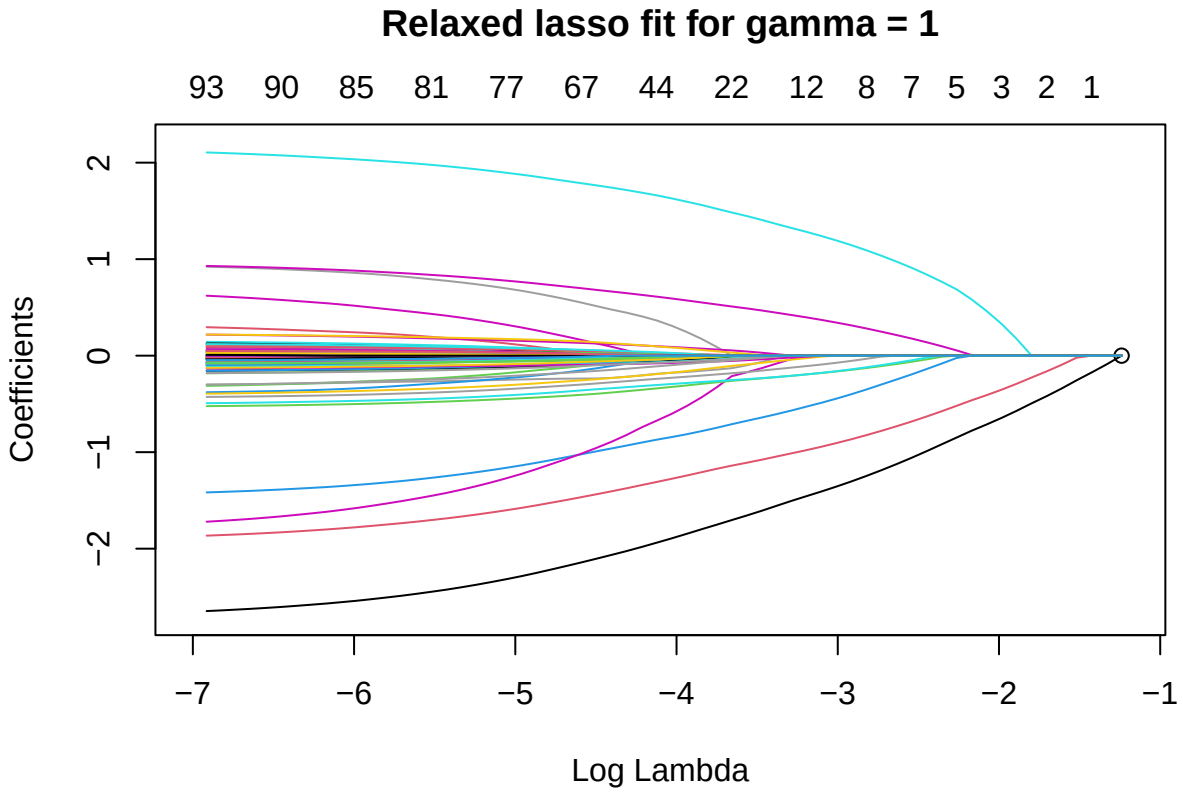


```
# Get an arbitrary set of coefficients for this example
lam = cox.fit$lambda[min(20,length(cox.fit$lambda))]
predict(cox.fit,lam=lam,gam=1)$beta
```

```
##          X4         X5        X18        X19        X21        X22        X24
## -0.4445542  1.1931259 -0.9052081 -0.1634296 -0.1607677  0.3425005 -0.0734942
##         X25
## -1.3556991
```

Nested cross validation

Because the values for lambda and gamma informed by CV are specifically chosen to give a best fit, model fit statistics for the CV derived model will be biased. To address this one can perform a CV on the CV derived estimates, that is a nested cross validation as argued for in SRDM (Simon R, Radmacher MD, Dobbin K, McShane LM. Pitfalls in the Use of DNA Microarray Data for Diagnostic and Prognostic Classification. J Natl Cancer Inst (2003) 95 (1): 14-18. https://academic.oup.com/jnci/article/95/1/14/2520188). This is done here by the nested.glmnetr() function.

```
# A nested cross validation to evaluate a cross validation informed lasso model fit
# nested.cox.fit = nested.glmnetr(xs,NULL,y_,event,family="cox",track=1)
nested.cox.fit = suppressWarnings(nested.glmnetr(xs,NULL,y_,event,family="cox",track=0))
summary(nested.cox.fit)
```

```
##
```

```
##  Sample information including number of records, events,number of columns in
##  design (predictor, X) matrix, and df (rank) of design matrix:
##     family          n    nevents xs.columns      xs.df
##      "cox"     "1000"      "337"      "100"       "94"
##
##  Tuning parameters for lasso/rpart model:
##     folds_n        seed
##        "10" "325180121"
##
##  Nested Cross Validation averages for LASSO (1se and min), Relaxed LASSO, and gamma=0 LASSO :
##
##       deviance per event :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      5.3569      5.2830      5.3446      5.3103      5.3808      5.3726
##
##       number of nonzero model terms :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##        26.1        50.7        21.9        36.7         7.0        11.7
##
##       linear calibration coefficient :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      1.3028      1.0918      1.2588      1.0469      1.0898      0.9771
##
##       agreement (concordance):
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      0.9193      0.9206      0.9194      0.9198      0.9137      0.9143
##
##  Naive agreement for cross validation informed lasso model :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      0.9253      0.9315      0.9253      0.9313      0.9288      0.9343
```

```
#names(nested.cox.fit)
```

Before providing analysis results the output first reports sample size and since this is for a Cox regression, the number of events, followed by the number of predictors and the df (degrees of freedom) of the design matrix, as well as some information on "Tuning parameters" which reflect the earlier work to compare the lasso model with stepwise procedures as described in JWHT (James, Witten, Hastie and Tibshirani, An Introduction to Statistical Learning with applications in R, 2nd ed., Springer, New York, 2021). In general we have found in practice that the lasso does better and so we do not present results here. (The tuned stepwise fits also take a long to run, part of the earlier motivation for the lasso model development.)

Next are the nested cross validation results. First are the per record (or per event in case of the Cox model) log-likelihoods which reflect the amount of information in each observation. Since we are not using large sample theory to base inferences we feel the per record are more intuitive, and they allow comparisons between datasets with unequal sample sizes. Next are the average number of model terms which reflect the complexity of the different models, even if in a naive sense, followed by the agreement statistics, here concordance, These nested cross validated concordances should be essentially unbiased for the given design, unlike the naive concordances where the same data are used to derive the model and calculate the concordances (see SRDM).

In addition to evaluating the CV informed relaxed lasso model using another layer of CV, the nested.glmnetr() function also runs cv.glmnetr() based upon the whole data set. Here we see, not unexpectedly, that the concordances estimated from the nested CV are slightly smaller than the concordances naively calculated using the original dataset. Depending on the data the nested CV and naive agreement measures, here concordance, can be very similar or disparate.

Following JWHT we provide information on the minimizing lasso models as well as the "1SE" models,

which are near to the minimizing lasso model fits, but of simpler nature. We though focus on the minimizing lasso fits recognizing that relaxed lasso and fully relaxed lasso fits generally provide models of simpler form while still optimizing a fit.

A summary for the CV fit can be produced by using the summary() function directly on a nested.glmnetr() output using the option *cvfit=TRUE*. Else one can also extract the CV fit by extracting the object$cv.glmnet.fit, where object is the output object obtained when running nested.glmnetr(). The plot() and predict() functions can be applied direclty to a nested.glmnetr() object without the *cvfit* option for futher evaluation or calculations for the CV model fit.

```
# Summary of the CV fit from a nested CV output
summary(nested.cox.fit, cvfit=TRUE)
```

```
##
##   The relaxed minimum is obtained for lambda = 0.01950023 , index = 30 and gamma = 0.75
##   with df (number of non-zero terms) = 36, average deviance = 5.407983 and beta =
##            X4            X5            X8           X14           X15
## -9.298178e-01  1.711721e+00 -1.885456e-01 -7.454261e-01  1.190965e-14
##           X16           X17           X18           X19           X20
##  4.167433e-01 -1.302087e-15 -1.362967e+00 -3.625131e-01  2.922682e-02
##           X21           X22           X23           X24           X25
## -3.182677e-01  6.426824e-01 -2.134638e-01 -2.569618e-01 -2.012441e+00
##           X29           X30           X34           X35           X38
##  2.389336e-02  2.816325e-02  2.602897e-02  2.119959e-02  1.064479e-01
##           X39           X43           X49           X50           X58
##  2.293174e-02 -5.218288e-02 -7.819052e-02  3.211853e-02  3.392804e-02
##           X59           X60           X61           X62           X67
## -3.948692e-02 -4.923017e-02 -7.808802e-02 -8.469542e-02  1.799604e-02
##           X71           X72           X77           X87           X96
## -4.703098e-02 -1.302645e-01  4.472362e-02  1.123160e-01 -8.482900e-02
##           X99
## -2.308179e-02
##
##   The fully relaxed (gamma=0) minimum is obtained for lambda = 0.03104988 and index = 25
##   with df (number of non-zero terms) = 17, average deviance = 5.445049 and beta =
##         X4         X5         X8        X14        X18        X19        X21
## -1.0991234  1.9492327 -0.4643855 -0.6377441 -1.5308945 -0.4075762 -0.3687472
##        X22        X23        X24        X25        X38        X61        X62
##  0.7625238 -0.2999622 -0.3593211 -2.2920690  0.1993155 -0.1477441 -0.1403819
##        X72        X87        X96
## -0.1784984  0.1874497 -0.1588458
##
##   The UNrelaxed (gamma=1) minimum is obtained for lambda = 0.01475121 and index = 33
##   with df (number of non-zero terms) = 48, average deviance = 5.415692
##
##
##   Order coefficients entered into the lasso model (1st to last):
##  [1] "X25" "X18" "X5"  "X22" "X4"  "X19" "X21" "X24" "X23" "X62" "X14" "X38"
## [13] "X8"  "X72" "X96" "X61" "X87" "X59" "X35" "X49" "X50" "X16" "X34" "X71"
## [25] "X77" "X30" "X39" "X43" "X58" "X60" "X20" "X29" "X67" "X99" "X12" "X28"
## [37] "X84" "X98" "X51" "X55" "X88" "X6"  "X11" "X33" "X78" "X40" "X66" "X68"
```

Observe, the summary here is slightly different than obtained above running cv.glmnetr(). This is because the model is derived using a new call (instance) of the cv.glmnetr() function, and each CV uses by default a new random partitioning of the data.

```
# Plot CV deviances from a nested CV output
plot(nested.cox.fit)
```

```
##  min CV average deviance (max log likelihood) for
##    relaxed at log(lambda)  = -3.937, gamma.min = 0.75, df = 36
##    fully relaxed at log(lambda)   = -3.472, df = 18
##    fully penalized at log(lambda) = -4.216, df = 48
```

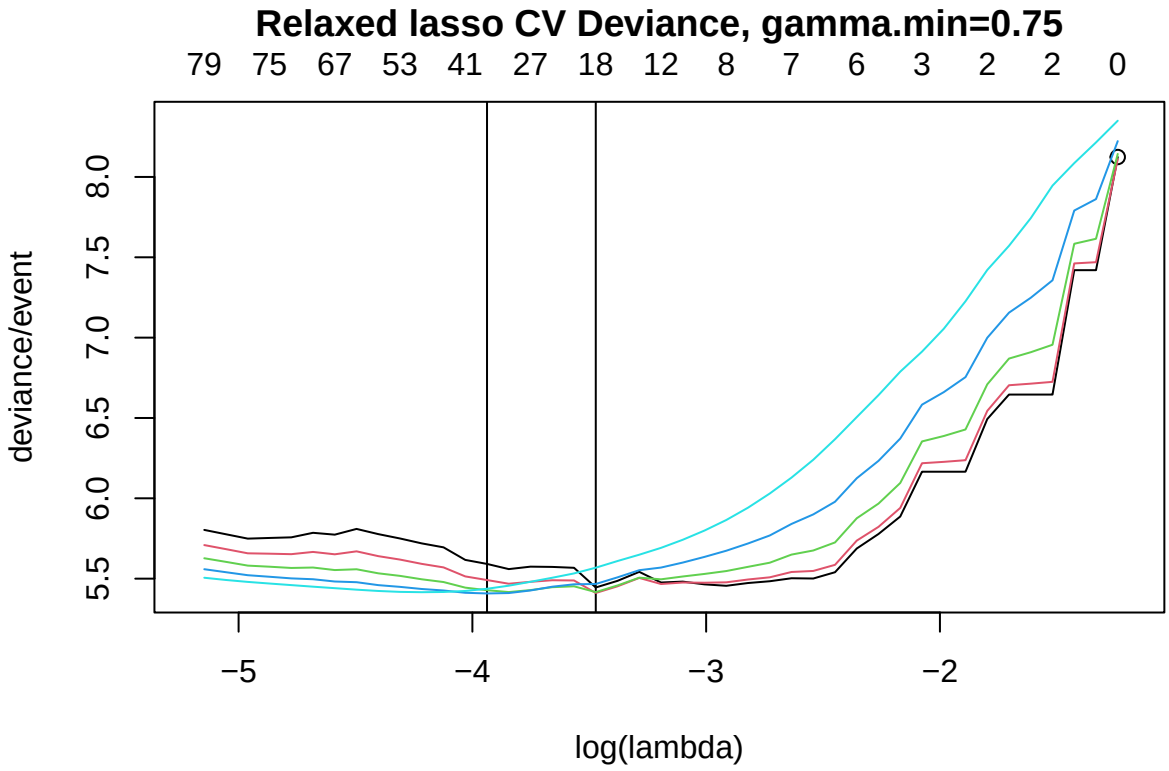

and

```
# Plot coefficients from a nested CV output
plot(nested.cox.fit, coefs=TRUE)
```

```
##  min CV average deviance (max log likelihood)
##    at log(lambda.min) = -3.937, gamma.min = 0.75, df = 36
```

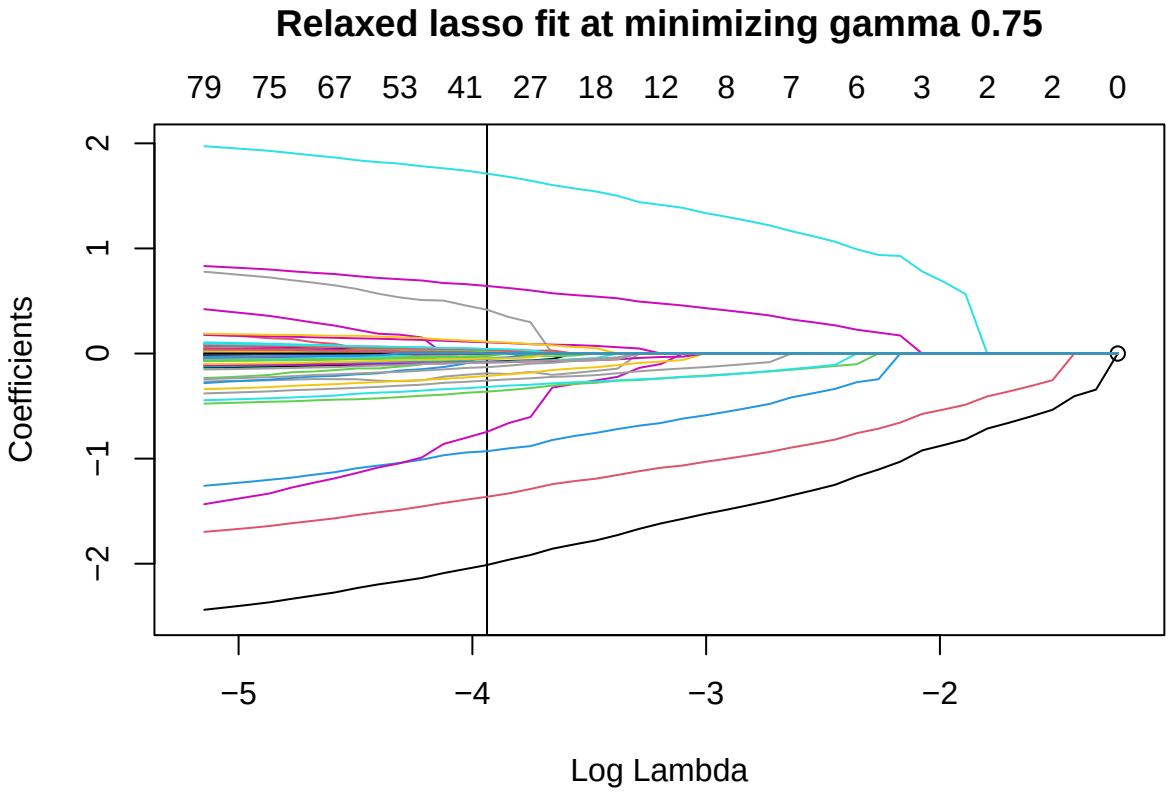

Summarizing, the summary() function with the *cvfit*=TRUE option as well as the plot() and predict() functions for a nested.glmnetr() object are then essentially the same as those for a cv.glmnetr() output object. The summary() function without the *cvfit*=TRUE option, though, regards the evaluation of the cv.glmnetr() fit and is different.

```
# ... or use the predict function on the CV fit embedded in the nested CV output
predict(nested.cox.fit)$beta
```

```
##
##  (lambda, gamma) pair minimizing CV average deviance is used

##            X4            X5            X8           X14           X15
## -9.298178e-01  1.711721e+00 -1.885456e-01 -7.454261e-01  1.190965e-14
##           X16           X17           X18           X19           X20
##  4.167433e-01 -1.302087e-15 -1.362967e+00 -3.625131e-01  2.922682e-02
##           X21           X22           X23           X24           X25
## -3.182677e-01  6.426824e-01 -2.134638e-01 -2.569618e-01 -2.012441e+00
##           X29           X30           X34           X35           X38
##  2.389336e-02  2.816325e-02  2.602897e-02  2.119959e-02  1.064479e-01
##           X39           X43           X49           X50           X58
##  2.293174e-02 -5.218288e-02 -7.819052e-02  3.211853e-02  3.392804e-02
##           X59           X60           X61           X62           X67
## -3.948692e-02 -4.923017e-02 -7.808802e-02 -8.469542e-02  1.799604e-02
##           X71           X72           X77           X87           X96
## -4.703098e-02 -1.302645e-01  4.472362e-02  1.123160e-01 -8.482900e-02
##           X99
## -2.308179e-02
```

Again, the plots and summary outputs from the nested.glmnetr() output are sightly different from what we saw above when summarizing the cv.glmnetr() output due to random data partitions for the CV folds.

More example data and relaxed lasso fits

The glmnetr.simdata() can be used to obtain example data not only survival analyses but also for linear models and logistic models. The glmnetr.simdata() output object list contains not only xs for the predictor matrix, yt for time to event or censoring and event for event indication but also y_ for a normally distributed random variable for the linear model setting and yb for the logistic model setting. Below we show examples extracting and analyzing simulated data and for the linear model and logistic model structures.

```
#=====================================================
# use the same simulated data output object from above, that is from the call
# simdata=glmnetr.simdata(nrows=1000, ncols=100, beta=NULL)
#
# extract linear regression model data
# xs = simdata$xs      # just as a comment as we did this above
yg = simdata$y_        # vector of Gaussian (normal) outcomes
# run a linear regression lasso model
cv.lin.fit = suppressWarnings(cv.glmnetr(xs,NULL,yg,NULL,family="gaussian",track=0))
summary(cv.lin.fit)
```

```
##
##   The relaxed minimum is obtained for lambda = 0.03319184 , index = 44 and gamma = 0.5
##   with df (number of non-zero terms) = 47, average deviance = 1.114769 and beta =
##            X4            X5            X8           X10           X11
##  1.161532e+00 -1.549189e+00  3.223005e-02 -3.006510e-01  2.265790e-01
##           X12           X14           X16           X17           X18
##  1.196422e-01  1.067253e+00 -8.194970e-01  7.280095e-14  1.269236e+00
##           X19           X20           X21           X22           X23
##  3.782328e-01 -1.159691e-01  4.013828e-01 -5.696146e-01  3.107055e-01
##           X24           X25           X31           X34           X38
##  3.021748e-01  1.820749e+00 -2.789714e-02 -4.220347e-02 -1.858647e-02
##           X39           X42           X43           X44           X49
## -3.955076e-02 -5.575306e-02  7.198290e-02 -3.456562e-02  3.776110e-02
##           X50           X58           X59           X60           X61
## -5.184501e-02  2.372919e-02  3.930152e-02  2.669017e-02  3.428079e-02
##           X62           X63           X67           X69           X72
##  4.135598e-02  2.856378e-02 -1.668165e-02  3.939001e-02  4.305880e-02
##           X74           X78           X79           X83           X85
##  4.264499e-02  1.403872e-02 -4.858182e-02  3.873856e-02 -3.775158e-02
##           X87           X89           X91           X93           X96
## -4.837261e-02  4.738807e-02  4.623275e-02  5.977543e-02  2.214418e-02
##           X97          X100
## -1.800447e-02  1.903081e-02
##
##   The fully relaxed (gamma=0) minimum is obtained for lambda = 0.05285079 and index = 39
##   with df (number of non-zero terms) = 29, average deviance = 1.129838 and beta =
##          X4          X5         X10         X11         X12         X14
##  1.20081248 -1.56120791 -0.37574751  0.26079629  0.17613590  1.22237202
##         X16         X18         X19         X20         X21         X22
## -0.96787628  1.28082345  0.39684005 -0.13359011  0.42388504 -0.58702071
```

```
##         X23         X24         X25         X42         X43         X49
##  0.31994642  0.31143832  1.84143659 -0.06802143  0.08228061  0.04727775
##         X50         X59         X61         X62         X72         X79
## -0.06763110  0.05154905  0.05362791  0.05577944  0.04777248 -0.06587503
##         X85         X87         X89         X91         X93
## -0.05354526 -0.06693239  0.05407041  0.06285454  0.07167155
##
##   The UNrelaxed (gamma=1) minimum is obtained for lambda = 0.01899359 and index = 50
##   with df (number of non-zero terms) = 68, average deviance = 1.11494
##
##
##   Order coefficients entered into the lasso model (1st to last):
##  [1] "X25"  "X18"  "X5"   "X4"   "X22"  "X21"  "X19"  "X23"  "X24"  "X14"
## [11] "X20"  "X7"   "X11"  "X16"  "X10"  "X43"  "X79"  "X12"  "X50"  "X93"
## [21] "X42"  "X62"  "X87"  "X89"  "X91"  "X49"  "X72"  "X59"  "X61"  "X85"
## [31] "X39"  "X44"  "X8"   "X34"  "X74"  "X83"  "X69"  "X31"  "X38"  "X63"
## [41] "X96"  "X58"  "X60"  "X67"  "X78"  "X97"  "X100" "X51"  "X75"  "X6"
## [51] "X48"  "X80"  "X82"  "X36"  "X57"  "X26"  "X52"  "X35"  "X45"  "X54"
## [61] "X56"  "X65"  "X71"  "X73"  "X90"  "X94"  "X46"  "X70"
```

```
# plot(cv.lin.fit, coefs=TRUE)
#
# extract logistic regression model data
# xs = simdata$xs      # just as a comment as we did this above
yb = simdata$yb        # vector of binomial (0 or 1) outcomes
# run a logistic regression lasso model
cv.bin.fit = suppressWarnings(cv.glmnetr(xs,NULL,yb,NULL,family="binomial",track=0))
summary(cv.bin.fit)
```

```
##
##   The relaxed minimum is obtained for lambda = 0.03486156 , index = 22 and gamma = 0
##   with df (number of non-zero terms) = 9, average deviance = 0.611035 and beta =
##         X4         X5        X18        X19        X21        X22        X23
##  1.8117021 -2.7690157  1.9554297  0.5723943  0.7843262 -0.7148489  0.3575476
##        X24        X25
##  0.4489259  2.6955189
##
##   The fully relaxed (gamma=0) minimum is obtained for lambda = 0.03486156 and index = 22
##   with df (number of non-zero terms) = 9, average deviance = 0.611035 and beta =
##         X4         X5        X18        X19        X21        X22        X23
##  1.8117021 -2.7690157  1.9554297  0.5723943  0.7843262 -0.7148489  0.3575476
##        X24        X25
##  0.4489259  2.6955189
##
##   The UNrelaxed (gamma=1) minimum is obtained for lambda = 0.00653242 and index = 40
##   with df (number of non-zero terms) = 53, average deviance = 0.618209
##
##
##   Order coefficients entered into the lasso model (1st to last):
##  [1] "X25"  "X18"  "X5"   "X4"   "X21"  "X22"  "X19"  "X24"  "X23"  "X20"
## [11] "X49"  "X11"  "X39"  "X68"  "X7"   "X10"  "X14"  "X47"  "X35"  "X50"
## [21] "X69"  "X37"  "X42"  "X61"  "X79"  "X91"  "X54"  "X40"  "X63"  "X93"
## [31] "X32"  "X72"  "X82"  "X16"  "X31"  "X92"  "X100" "X34"  "X51"  "X85"
## [41] "X41"  "X67"  "X81"  "X30"  "X44"  "X66"  "X74"  "X84"  "X9"   "X52"
```

```
## [51] "X55"  "X36"  "X38"
```

```
# plot(cv.bin.fit, coefs=TRUE)
```

Further model assessment

One can also fit a spline to the predicteds obtained form the predict functions. This may help to understand nonlinearities in the predicteds, but may also give inflated hazard ratios.

```
# Get predicteds from CV relaxed lasso model embedded in nested CV outputs & Plot
xb.hat = predict( object=cv.cox.fit , xs_new=xs, lam=NULL, gam=NULL, comment=FALSE)
# describe the distribution of xb.hat
round(1000*quantile(xb.hat,c(0.01,0.05,0.1,0.25,0.5,0.75,0.90,0.95,0.99)))/1000
```

```
##     1%     5%    10%    25%    50%    75%    90%    95%    99%
## -5.912 -4.444 -3.633 -1.904  0.039  1.915  3.454  4.615  6.414
```

```
# Fit a spline to xb.hat uisng coxph, and plot
fit1 = coxph(Surv(y_, event) ~ pspline(xb.hat))
termplot(fit1,term=1,se=TRUE)
```

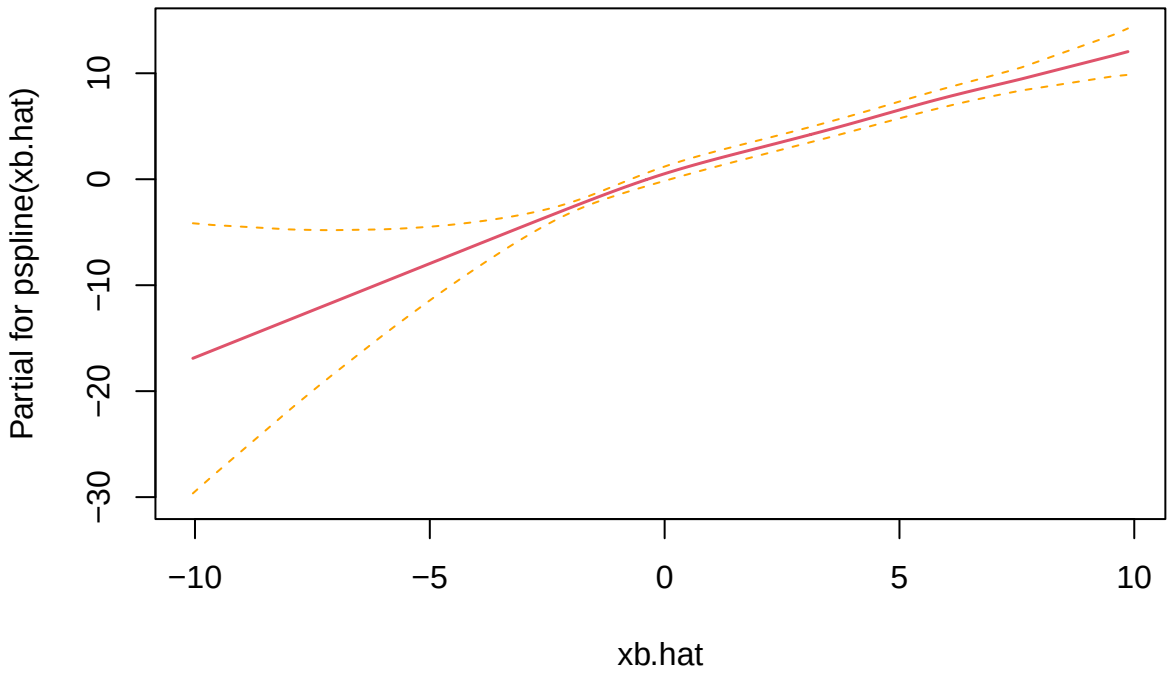

From this spline fit we see the predicteds are approximately linear with the log hazard ratio.