Using stepreg

Walter K. Kremers, Mayo Clinic, Rochester MN

10 December 2022

The Package

The stepreg() and cv.stepreg() funcitons in the *glmnetr* package were written for convenience and stability as opposed to speed or broad applicability. When fitting lasso models we wanted to compare these to standard stepwise regression models. Keeping a more modern approach we tune by either number of terms included in the model (James, Witten, Hastie and Tibshirani, An Introduction to Statistical Learning with applications in R, 2nd ed., Springer, New York, 2021) or by the p critical value for model inclusion, as this too is a common tuning parameter when fitting stepwise models.

When fitting lasso models we often use one-hot coding for predictor factors when setting up the design matrix. This allows lasso to identify and add to the model a term for any one group that might be particularly different from the others. By the penalty lasso stabilizes the model coefficients and keeps them form going to infinite, while ridge will generally uniquely identify coefficients despite any strict collinearities.

Before writing this program we tried different available packages to fit stepwise models for the Cox repregreison framework but all we tried had difficulties with numerical stability for the large and wide clinical datasets we were working with, and which involved one-hot coding. There may well be a package that would be stable for the data we were analyzing but we decided to write this small function to be able to tune for stability.

This program is slow but our goal was not for routine usage but to use the stepwise procedure on occasion as a reference for the lasso models. For many clinical datasets the lasso clearly outperformed the stepwise procedure, and ran much faster. For many simulated data sets with simplified covariance structures, i.e. independence of the underlying predictors, the lasso did not appear to do much better than the stepwise procedure tuned by number of model terms or p.

Data requirements

The data requirements for stepreg() and cv.stepreg() are similar to those of cv.glmnetr() and refer to the *Using glmneter* vignette for a description.

An example dataset

To demonstrate usage of *cv.stepreg* we first generate a data set for analysis, run an analysis and evaluate. Following the *Using glmnetr* vignette, the code

```
# Simulate data for use in an example survival model fit
# 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 *cv.stepreg* (and *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 stepwise 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() function.

```
# Fit a stepwise regression model informed by cross validation
cv.stepwise.fit=suppressWarnings(cv.stepreg(xs,NULL,y_,event,family="cox",folds_n=5,steps_n=30))
```

Note, in the derivation of the stepwise regression models, individual coefficients may be unstable even when the model may be stable which elicits warning messages. Thus we "wrapped" the call to cv.stepreg() within the suppressWarnings() function to suppress excessive warning messages in this vignette. The first term in the call to cv.stepreg(), 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. | To summarizing the model fit and inspect the coefficient estimates we use the summary() function.

```
# summarize model fit ...
summary(cv.stepwise.fit)
```

```
##
##  CV best df = 20, CV best p enter = 0.05 for 24 predictors
##      in the full data model, from 100 candidate predictors
##
##   df loglik.null    loglik     pvalue concordance         std        X4
## 1 20   -2141.057 -1581.991 0.02620091   0.9309904 0.005033834 -1.332723
## 2 24   -2141.057 -1573.472 0.03365787   0.9314918 0.004936426 -1.304863
##         X5        X6         X8       X14      X16       X18        X19
## 1 1.963749 0.6193282  0.0000000 -1.834483 1.181275 -1.646132 -0.5048882
## 2 2.042968 0.5496746 -0.3118365 -1.773074 1.055766 -1.701096 -0.5159846
##         X20        X21       X22        X23        X24       X25       X38
## 1 0.1614991 -0.3979934 0.8094201 -0.3670110 -0.3721133 -2.417468 0.1903912
## 2 0.1238821 -0.4051913 0.8221877 -0.3839914 -0.3826572 -2.476905 0.2062551
##          X43        X49        X60        X61        X62        X71        X72
## 1  0.0000000 -0.1644987 -0.1234921  0.0000000 -0.1281786  0.0000000 -0.2350263
## 2 -0.1266882 -0.1714386 -0.1165997 -0.1348102 -0.1354011 -0.1288603 -0.2352343
##         X87        X96
## 1 0.2310910 -0.1675236
## 2 0.2306497 -0.1659381
```

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 cross validation informed lasso and stepwise fit
y_ = simdata$y_
nested.cox.fit = suppressWarnings(nested.glmnetr(xs,NULL,y_,NULL,family="gaussian",
                                  dolasso=1,dostep=1,doaic=1,folds_n=3,steps_n=30))
#nested.cox.fit= suppressWarnings(nested.glmnetr(xs,NULL,y_,event,family="cox",
#                                  dolasso=1,dostep=1,doaic=1,folds_n=3,steps_n=30))
summary(nested.cox.fit)
```

```
##
##  Sample information including number of records, number of columns in
##  design (predictor, X) matrix, and df (rank) of design matrix:
##     family          n xs.columns      xs.df
## "gaussian"     "1000"      "100"       "94"
```

```
##
##  Tuning parameters for models:
##     steps_n     folds_n      method     dolasso       doaic      dostep
##        "30"         "3"    "loglik"         "1"         "1"         "1"
##        seed
## "171989386"
##
##  Nested Cross Validation averages for LASSO (1se and min), Relaxed LASSO, and gamma=0 LASSO :
##
##       deviance per record :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      1.1885      1.1069      1.1821      1.1169      1.1823      1.1895
##
##       number of nonzero model terms :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##       25.33       53.67       22.67       39.67       12.00       21.33
##
##       linear calibration coefficient :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      1.0710      1.0300      1.0563      1.0253      0.9956      0.9873
##
##       agreement (concordance for Cox and binomial, r-square for guassian):
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      0.8572      0.8644      0.8574      0.8632      0.8542      0.8537
##
##  Naive agreement for cross validation tuned lasso model :
##   lasso.1se   lasso.min  lasso.1seR  lasso.minR lasso.1seR0 lasso.minR0
##      0.8664      0.8784      0.8664      0.8784      0.8764      0.8829
##
##  Nested Cross Validation stepwise regression model (df):
##       Average deviance : 1.1195
##       Average model df : 14
##       R-square         : 0.8627
##  Naive R-square based upon the same (all) data as model derivation (df): 0.8703
##
##  Nested Cross Validation stepwise regression model (p):
##       Average deviance : 1.1275
##       Average model p  : 0.013
##       Average model df : 14.33
##       R-square         : 0.8617
##  Naive R-square based upon the same (all) data as model derivation (p): 0.8717
##
##  Cross Validation results for stepwise regression model: (AIC)
##       Average deviance : 1.1556
##       Average model df : 29
##       Concordance      : 0.8586
##  Naive R-square based upon the same (all) data as model derivation (AIC) : 0.8782
```

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

For this example we use only 3 folds, instead of 5 or 10 like we would do in practice, to get reasonble run times as this is just for the purpose of demostration.
| 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" 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, Springer, New York, 2021). In general we have found in practice that the lasso performs better.
| 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 model fits using another layer of CV, the nested.glmnetr() function does the CV fits 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 R-square, can be very similar or disparate.
| A summary for the CV fit can be produced by extracting the object$cv.stepreg.fit, where object is the output object obtained when running nested.glmnetr(), and then using the summary() function.

```
# Summary of the CV informed stepwise model fit
summary(nested.cox.fit$cv.stepreg.fit)
```

```
##
##  CV best df = 14, CV best p enter = 0.02 for 16 predictors
##      in the full data model, from 100 candidate predictors
##
##   df loglik.null    loglik     pvalue   rsquare rsquareadj        Int       X4
## 1 14   -2464.276 -1443.094 0.01248035 0.8702784  0.8684347 -0.1699559 1.201499
## 2 16   -2464.276 -1437.412 0.01657311 0.8717442  0.8696566 -0.1777801 1.199209
##          X5        X10       X11      X14      X17      X18       X19
## 1 -1.613359 -0.4090516 0.0000000 1.268976 1.062544 1.284108 0.3921502
## 2 -1.598965 -0.3741456 0.2660345 1.277813 1.022508 1.280104 0.3974841
##          X20       X21        X22       X23       X24      X25         X50
## 1 -0.1360695 0.4167878 -0.5976320 0.3104475 0.3198516 1.834711 -0.07914720
## 2 -0.1362363 0.4231258 -0.5952538 0.3127745 0.3210119 1.836684 -0.07877755
##          X93
## 1 0.00000000
## 2 0.08334172
```

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