

APTS Statistical Computing 2022/23: Practical Lab 1 (Wednesday)

Here are some practical problems which aim to explore and reinforce some of the course material. They all use R. This is just for convenience: many real statistical numerical analysis tasks are approached with a mixture of compiled code, using a statically typed language (like C, C++, Fortran, Java, Scala, . . .), together with code in a high level language such as R, Python or Matlab. But for exploratory and pedagogic purposes, high level languages are often convenient. Some languages (such as Scala and Julia) attempt avoid the “two language problem” by being both high level and easy to develop in interactively, as well as fast and efficient.

Several problems here use simulated data: when developing statistical modelling code, it is often best to start out with data where you know what the truth is (and can generate further replicates). Do not consult the solutions¹ until you’ve made good attempts.

1. Download the example code extracted from the course notes, from the course website:

```
https://darrenjw.github.io/work/teaching/aps-sc/
```

Load this code into R-Studio (or your favourite R IDE), and execute some code blocks to make sure that you can reproduce some of the examples from the lecture notes. Some of the code examples from the notes will be useful for solving some of the exercises.

2. *Principal components analysis (PCA)* of a multivariate data set was traditionally based on the eigen-decomposition of the sample covariance matrix of the data. The matrix of eigenvectors can be used to rotate the (centred) data observations to a set of uncorrelated random quantities ordered by decreasing variance. These rotated data are often known as the *scores*. We can write a small function to implement this as follows.

```
pcScoresEig = function(X) {  
  eig = eigen(var(X), symmetric=TRUE)  
  Xc = sweep(as.matrix(X), 2, colMeans(X))  
  Xc %*% eig$vector  
}
```

We can test it on the infamous iris data, as follows.

```
Xi = iris[,-5]  
scores = pcScoresEig(Xi)  
plot(scores[,1], scores[,2], col=iris[,5], pch=19)
```

This is essentially how the `princomp` function in R is implemented, and we can verify this.

```
head(scores, 3)  
head(princomp(Xi)$scores, 3)
```

¹Solutions will be made available from the course website before the end of the session.

- (a) It turns out that the singular value decomposition of the (centred) data matrix can be used to construct the scores directly as \mathbf{UD} . Think about why this is true, and write an R function, `pcScoresSvd` to implement this. Test it on the iris data, and don't worry about a sign flip.
- (b) This SVD-based method is (possibly) a bit cheaper, and much more numerically stable than the eigendecomposition method, and works for data that is not of full rank, including wide data ($p > n$). It also gives some additional insight into what the SVD "means". This is essentially how the `prcomp` function in R is implemented (which is almost always preferred to the `princomp` function). Compare your function with this.
- (c) Simulate some random (eg.) $5,000 \times 1,000$ test data, and time your two implementations. The relative timings are very sensitive to precise details of underpinning BLAS and LAPACK implementations. I typically find that the SVD method is slightly faster, but that isn't really the point — the point is that there isn't much difference, but the SVD method is more numerically stable.
- (d) Think about how to use the SVD to compute the variances or standard deviations of the principal components. The variances would be given by the diagonal of the \mathbf{D} matrix for the eigendecomposition approach, but this doesn't quite work for the SVD approach. Modify your function to compute them. Check it against `prcomp` for the iris data to make sure you've done it correctly.

3. For a linear regression model,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$$

we know that the quadratic loss $L_0(\boldsymbol{\beta}) = \|\boldsymbol{\epsilon}\|^2 = \boldsymbol{\epsilon}^T \boldsymbol{\epsilon}$ is minimised wrt $\boldsymbol{\beta}$ when $\boldsymbol{\beta}$ is a solution to the normal equations,

$$\mathbf{X}^T \mathbf{X} \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}.$$

In *ridge regression*, the slightly modified quadratic loss function $L_\lambda(\boldsymbol{\beta}) = \|\boldsymbol{\epsilon}\|^2 + \lambda \|\boldsymbol{\beta}\|^2$ is used, for some ridge penalty $\lambda > 0$, which encourages shrinkage of the regression coefficients towards zero.

- (a) Show that for a given fixed $\lambda > 0$, the loss $L_\lambda(\boldsymbol{\beta})$ is minimised when $\boldsymbol{\beta}$ is a solution to

$$(\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I}) \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}.$$

- (b) Starting from the singular value decomposition, $\mathbf{X} = \mathbf{UDV}^T$, show that the optimal $\hat{\boldsymbol{\beta}}_\lambda$ can be written as

$$\hat{\boldsymbol{\beta}}_\lambda = \mathbf{VD}_\lambda \mathbf{U}^T \mathbf{y},$$

where \mathbf{D}_λ is a diagonal matrix with entries $d_i^\lambda = d_i / (d_i^2 + \lambda)$. Note that this means $\hat{\boldsymbol{\beta}}_\lambda$ can be computed for as many different λ as desired, all for the cost of one single expensive SVD operation.

- (c) In practice, both the data, \mathbf{y} , and the covariate matrix \mathbf{X} are centred before ridge regression is applied, since then the model can be fit without an intercept, and typically you would not want to shrink the intercept. Write a function,

```
ridge(y, X, lambda)
```

which expects an n -vector \mathbf{y} , an $n \times p$ matrix \mathbf{X} , and a q -vector of λ values where the ridge solution is required. The function should return a $p \times q$ matrix of ridge regression parameters, with each column representing a solution for a given λ .

- (d) For the `trees` dataset, regress volume on the other two variables for a range of shrinkage parameters.

```
ridge(trees[,3], trees[,1:2], c(0, exp(0:5)))
```

Ensure that your solution matches up with that of `lm` in the case $\lambda = 0$.

4. This question revisits the third example in section 1.1 of the notes.

(a) Run the following code that reproduces the good and bad fits produced by calls to `lm` in the example.

```
## First, simulate the data
set.seed(1)
x <- sort(runif(100)) + 100
y <- .2*(x-100 -.5) + (x-100 -.5)^2 + rnorm(100)*.1

## lm manages to estimate the regression
plot(x, y)
b <- lm(y~x+I(x^2))
lines(x, fitted(b), lwd=2)

## direct solution of the normal equations doesn't work
X <- model.matrix(b)
beta.hat <- solve(t(X)%*%X, t(X)%*%y)

## What if the x-values were even further away from zero?
## In theory, this should still be able to represent the quadratic function.
x1 <- x+1000
plot(x, y)
b1 <- lm(y~x1+I(x1^2))
lines(x, fitted(b1), lwd=2)
## lm fails without warning!
```

- (b) Examine the condition numbers of the model matrices in the two cases (`lm` computes the fit using the QR decomposition approach, not by direct solution of the normal equations). Why was the second `lm` fit so bad?
- (c) Plot the second and third columns of the model matrix against each other for the two cases, and use `cor` to examine their correlation. Why are the condition numbers so high here?
- (d) Since the linear model says simply that the expected value vector $E(\mathbf{y})$ lies in the space spanned by the columns of \mathbf{X} , one possibility is to attempt to arrive at a better conditioned \mathbf{X} by linear rescaling and/or recombination of its columns. This is always equivalent to a linear reparameterization. Try this on the model matrix of the model that causes `lm` to fail. In particular, for each column (except the intercept column) subtract the column mean. Then divide each column (except the intercept column) by its standard deviation. Find the condition number of the new model matrix. Fit the model with this model matrix using something like `lm(y ~ Xs - 1)` where X_s is the re-scaled model matrix. Produce a plot that confirms that the resulting fit is sensible now.
- (e) An alternative fix is to subtract the mean x value from the original x vector before fitting. Try this and see what happens to the condition number and column correlations of the model matrix now.

- (f) If we want good condition numbers then the best thing would be have a model matrix made up of columns from an orthogonal matrix. *Orthogonal polynomials* provide a way of achieving this. Try fitting with `lm(y~poly(x,2))` to see these in action. Look at the correlation between the model matrix columns now, and the condition number.
- (g) The model matrix produced in the last part is not quite (column) orthogonal (meaning orthonormal), but to correct this we could rescale and use a model matrix `X <- cbind(n^-.5,poly(x,2))` where `n` is the number of data. Without computing anything new, find a QR decomposition of this new model matrix? Show how the new model matrix can be used directly to find the fitted values in this case, without any need for an `lm` call. Check that it works.