Wishart

Paul Johnson <pauljohn@ku.edu> & James W. Stoutenborough <jstout@ku.edu>

June 10, 2013

1 Introduction.

Suppose you have drawn an observation from a Wishart distribution. What do you have?

You have a matrix! Its a covariance matrix. A square set of numbers, its symmetric

$$W = \begin{bmatrix} w_{11} & w_{12} & w_{13} & w_{14} \\ w_{21} & w_{22} & w_{23} & w_{24} \\ w_{31} & w_{32} & w_{33} & w_{44} \\ w_{41} & w_{42} & w_{43} & w_{45} \end{bmatrix}$$

It is also positive definite. Positive definite means that it would be a suitable covariance matrix for a multivariate normal distribution. Technically, if we multiply that by any vector x like this

$$x'Wx > 0$$

If we have any matrix–literally any matrix X–and calculate $X'X$, the result is a sum of squares and crossproducts matrix, which is symmetric and positive definite.

Thus one lazy person's way to create Wishart variables is to draw p-variate multivariate normal observations, stack them into an Nxp matrix (X), and then multiply $X'X$ to create a covariance matrix.

Run this in R, using the handy multivariate normal generator from the package mvtnorm:

```
library(mvtnorm)
meanV <- c(0,0,0,0)
varV <- 3*diag(4)
X <- rmvnorm(10, mean = meanV, sigma = varV)
W <- t(X) %*% X
W
```

```
           [,1]       [,2]        [,3]      [,4]
[1,] 28.6592382   6.291899  -0.6252022 -2.311278
[2,]  6.2918986  47.596049 -19.0211707 -5.709700
[3,] -0.6252022 -19.021171  20.2127947  6.125826
[4,] -2.3112780  -5.709700   6.1258257 36.618713
```

How about another Wishart? Re run, recalculate $X'X$?

```
              [,1]        [,2]       [,3]        [,4]
[1,]   15.7263306  -13.569627  -8.283313  -0.5315097
[2,]  -13.5696273   69.715530   9.415498 -25.3579642
[3,]   -8.2833134    9.415498  35.270566 -12.0242624
[4,]   -0.5315097  -25.357964 -12.024262  54.5802508
```

There's more where that came from:

```
            [,1]        [,2]        [,3]        [,4]
[1,] 21.982553   6.7582014  -3.6663810   7.1137074
[2,]  6.758201  41.7811616  -0.6594455  -2.7048963
[3,] -3.666381  -0.6594455  28.0277308   0.2415728
[4,]  7.113707  -2.7048963   0.2415728  42.3190939
```

Are you getting the idea? It is like any other probability model we have studied, except this one is different because the outcome is a whole matrix. Each draw produces a block of wholly new numbers. It is not a single number like the Normal distribution. Its not a vector like multivariate Normal. Instead, its a whole block. The outcomes are blocks of numbers. How many ways can I say it?

Is there any restriction on these values? Yes. They are sums of squares and cross products.

Just to be clear, we assume we are able to draw multivariate normal observations with a mean of 0 and a covariance matrix Σ. It has to be symmetric and positive definite, of course.

If we "stack" those p-variate normal draws into a matrix that is N rows and p columns wide, the result of $X'X$ follows a Wishart distribution. The number of rows in the matrix, which I called N here, plays the role of "degrees of freedom". The p x p matrix Σ can be any positive definite matrix you choose.

R has a function for drawing samples from any Wishart distribution. It is called rWishart and in order to draw from it, we must specify the covariance matrix. Suppose the covariance matrix is Σ, "Sigma" and the number of degrees of freedom. So far as I can tell, however, it is not doing anything fancy. It fills up a matrix of X, row-by-row, of p-variate normals, and then calculates X'X .

2 Mathematical Description

Please review the writeup on the Chi-Square distribution. Recall a Chi-Square variable can be thought of as the result of sampling from a standard normal distribution and then squaring the elements and summing them. In vectors, think of x as a column vector. Then the sum of squares is:

$$x' \cdot x = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}$$

If you draw a set of observations from this multivariate Normal distribution, then when you multiply, you'd have

$$X'X$$

The number of rows in X is the number of random samples drawn, and it is called the "degrees of freedom." The number of columns is equal to the desired number of rows (= columns) of the Wishart output.

The Wishart distribution is the multivariate analog to the chi-square, in the sense that the Wishart describes the variation you would observe if you repeatedly drew samples and calculated $X'X$. It is related to the multivariate normal in the same way that the chi-square is related to the univariate normal.

The Wishart draw w_i might be thought of as a reflection of the PRECISION of an MVN distribution. If the covariance matrix of the MVN is called Σ, then its inverse, Σ^{-1} , is its precision.

3 The Probability Denisty

The Wishart distribution can be characterized by its probability density fuction, as follows.

Let W_i be a $m \times m$ symmetric matrix which a randomly drawn value according to a Wishart distribution. The PDF of W_i is:

$$prob(W_i|v,S) \propto |W_i|^{(v-m-1)/2} exp(-\frac{1}{2} trace(S \cdot W_i))$$

The "trace" is the sum of all diagonal elements.

By construction, W_i is "positive definite", so it can serve as a precision matrix in a MVN distribution. The parameter v is "degrees of freedom" and S is a "cross product" matrix. See Lancaster, p. 179.

4 Moments

The expected value of a Wishart is

$$E(W_i) = v\,S^{-1}$$

5 About the example code.

The multivariate Normal is used here, and I chose a particularly simple covariance matrix (mostly because I was lazy).

$$\Sigma_4 = 3 * \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The inverse of that very simple matrix is (obviously)

$$1/3 * \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

This plays the role of S^{-1} in the model.