

R Programming: Worksheet 5

By the end of today you should understand the `apply()` family of functions, as well as creating and manipulating multi-way arrays.

1. Warm Up

Use `apply()`, `tapply()`, `sapply()`, `mapply()` and `replicate()` to solve these problems.

- (a) Generate a list of 50 random datasets, each consisting of 10 independent t_5 distributed random variables. [The relevant function is `rt()`.]
- (b) Generate a list of length 20, where the i th entry in the list is the sequence of numbers $1, \dots, i$.
- (c) Generate the following random matrix `X`:

```
> set.seed(2014)
> X <- matrix(rexp(200), 20, 10)
```

Find the smallest entry in each column of `X`.

- (d) Look at the data frame `C02` (this is preloaded into R). How would you determine which columns are numeric? [Hint: there is a function called `is.numeric()`]

2. Death Penalty Data

Here is a famous data set which is used to illustrate *Simpson's paradox*. It records the races of the victim and defendants in various murder cases in Florida between 1976 and 1987, and whether or not the death penalty was imposed upon the killer. The data are presented as counts.

Victim	White		Victim	Black	
Defendant	White	Black	Defendant	White	Black
Yes	53	11	Yes	0	4
No	414	37	No	16	139

- (a) Create a 3-way array in R to hold these data, with the three dimensions representing the penalty, the defendant's race and the victim's race respectively.

You can include suitable dimension names on a table with the `dimnames()` function.

```
> house <- dget("housing.dat")
> dimnames(house)
```

This is just a list of names for each dimension.

- (b) Create a similar list of dimension names for your death penalty data, called (say) `listOfNames`. Set it to the table by running

```
> dimnames(tab) <- listOfNames
```

- (c) Obtain the marginal distribution (i.e. summing out over the other variables) of the penalty imposed against the defendant's race.
- (d) Use this to estimate the probability of imposing the death penalty conditional upon race of the defendant. What do you conclude?
- (e) Now do the same but only for white victims. What do you conclude? What about only for black victims?
- (f) * Can you explain how this is possible?

[You can also use the function `conditionTable()` in the excellent(!) package `rje` to do the bits in part (d) and (e).]

3. Central Limit Theorem

- (a) Write a function with arguments `n` and `k`, which generates `k` data sets each consisting of `n` independent uniform random variables on $(0,1)$, and returns their mean. Do this in **three** different ways, using (i) a loop; (ii) `replicate()`; (iii) the function `rowMeans()`. [The function `runif()` may be useful.]
- (b) Using `system.time()`, time your functions for a few different values of `n` and `k`. Which is the fastest? Which function is the easiest for you to read and understand?
- (c) Run your preferred function for $k = 10,000$ and each of $n = 1, 2, 3, 5, 10$; plot the results as a histogram in each case. (Use `hist()`, and play around with the `breaks` and `col` options.)
- (d) For the case $n = 5$, plot a normal density over your histogram with the same mean and variance as the random variables generated by your function: here I mean the *exact* mean and variance. The function `dnorm()` may be useful. [Hint: recall that the variance of a Uniform(0,1) random variable is $\frac{1}{12}$.]

4. Confidence Intervals

- (a) Write a function which, given a real vector `x`, calculates a 95% confidence interval for the mean, using the usual t -distribution approximation. [See `?qt`]
- (b) Write a second function with argument `n`, which generates 20 data sets each consisting of `n` standard normal random variables, and calculates a 95% confidence interval for each data set. The function should return a list containing two vectors of length 20 named `lower` and `upper`, consisting of the lower and upper ends of each interval.
- (c) Write a third function which obtains 20 confidence intervals as above, and then plots them. The plot should consist of 20 horizontal lines, one drawn below the other.
Plot these a few times and watch the lines dance around!
- (d) Improve your plot function in the following ways:
 - (i) make the true mean always be in the centre of the plot;
 - (ii) plot intervals which cover the true mean in a different colour.
- (e) Try changing the mean you use to generate the data, and see how this changes your plot.
- (f) * How might you use this function to explain the frequentist interpretation of confidence intervals to a student?