

Homework 5

Biology 607

10/5/2014

1 χ^2 - the function!

1.1 (5 points)

Write a simple function to calculate a chi-square test based on the assumption that you should have the same observed values across all categories. Have it return the value of χ^2 , the degrees of freedom, and the p-value. To test it, compare your function's output for the observed vector 1:10 versus R's `chisq.test` function.

1.2 (5 points)

Use your function to answer question 14 from Chapter 8 of W&S. Answer all parts of the question (some will not require R - sad, I know.)

2 Contingency Tables

In W&S 9.3, we are introduced to the χ^2 contingency test which seeks to evaluate association between categorical variables - such as whether a fish is infected or uninfected with a parasite v. whether or not it is eaten by birds. Using the formulae in your book write a function that will return the χ^2 , DF, and p-value for a contingency table. Use the following matrix as test data (from the book) as well as the output from R's own `chisq.test`. Show me that your function works.

```
##           Uninfected Infected Highly Infected
## Eaten           1         10             37
## Not Eaten       49         35             9
##
## Pearson's Chi-squared test
##
## data:  eaten
## X-squared = 69.76, df = 2, p-value = 7.124e-16
```

Some helpful functions for you - `nrow`, `ncol`, `rowSums`, `colSums`. Also, read the whole section before writing this. In writing the function, one thing that might help would be to write out what you're going to do, step by step, in comments. Each comment should be one step. Then fill in the function with code that implements what you have written in each comment. Heck, ponder adopting this as your standard workflow. It will save you a lot of time.

2.1 (5 points)

Show me that your function works by answering question 21 from Chapter 9 of W&S. Note, the default behavior of R's `chisq.test` on the same data set is to invoke Yates' continuity correction. So, careful when checking yourself.

3 Student's T and Power

Have a Guinness. Are you down? Now, there are some exact formulae to calculate the power of a t-test. Let's see how they compare to a simulation based approach. Let's compare the power via simulation for a sample with a sample size of 10, a mean of 2, and a standard deviation of 3 to the one calculated by conventional means. In R, to do this conventionally, we'd use the following for a 2-tailed one sample T-test.

```
power.t.test(n=10, delta=2, sd=3,
             alternative="two.sided", type="one.sample")

##
##      One-sample t test power calculation
##
##              n = 10
##             delta = 2
##              sd = 3
##      sig.level = 0.05
##             power = 0.4692
##      alternative = two.sided
```

How about simulation? Let's put it together using functions. This question will walk you through it.

3.1 T-Test! (5 points)

To start with, let's write a function for the calculation of p-values for our data and evaluate it. Write a function to conduct a one sample t-test. Run it against sample of 10 values drawn from a normal distribution with mean 2 and sd of 3. Check the results against R's `t.test` and show me that both agree.

3.2 In which we begin simulations (5 points)

Now we need to get the power of our t-test via simulation. We'll accomplish this in two steps. First, write a function that will give you a vector of p values from t-tests on simulated data. This function should accept the number of simulations, the sample size, an critical effect size (the delta from above), and a sample standard deviation as it's arguments at the very least. Use it to create a vector of 5000 simulated p values using a sample size of 10, an effect size of 2, and a sd of 3, as above.

3.3 Simulated Power! (5 points)

Now that you have a vector of simulated p-values, write a function that will take such a vector, and calculate it's power given an alpha. What is the power of your test, given an alpha of 0.05? How does it compare to the power from the exact calculation above?