

Mathematics 376 – Mathematical Statistics
Computer Lab Day 2 in R – Confidence Intervals
February 20, 2012

Goals

The goals of today’s lab are:

- To gain some more intuitive understanding of meaning of confidence intervals by developing a graphical demonstration along the lines of Figure 8 on page 414 of our text. The concept of confidence intervals is a notoriously “slippery” one. It is easy to go from the intuitively appealing “confidence level” $(1 - \alpha) \times 100\%$ to statements that seem to be saying the same thing, but are misleading at best, and completely meaningless at worst. To use this idea reliably to make inferences from real-world data, it is very important to understand *exactly* what it means.
- To gain some experience with user-developed R functions.

A Confidence Interval Demo

If we have Y_1, \dots, Y_n in the “large sample” case ($n \geq 30$ or so), then we have seen that an approximate $(1 - \alpha) \times 100\%$ confidence interval for the population mean μ is produced by the formula:

$$\left(\bar{Y} - z_{\alpha/2} \frac{S}{\sqrt{n}}, \quad \bar{Y} + z_{\alpha/2} \frac{S}{\sqrt{n}} \right)$$

But what does this really mean, and what is the right way to interpret the “confidence level” $(1 - \alpha) \times 100\%$?

The following R code will generate a collection of *one hundred* 95% confidence intervals based on samples of size $n = 50$ from a normal population with mean $\mu = 10$ and $\sigma = 2$. It produces two things:

- (a) a graphic window display – each interval is drawn with a horizontal line, and the location of the population mean is marked by the vertical line
- (b) the number of intervals containing μ , printed in the R console window as the output of the function.

```
confdemo <- function()
{
  endpts <- matrix(1:200,nrow=100,byrow=T)
  nin <- 0
  for (i in 1:100)
  {
    ysamp <- rnorm(50,10,2)
```

```

ybar <- mean(ysamp)
S <- sd(ysamp)
endpts[i,1] <- ybar - qnorm(.975,0,1)*S/sqrt(50)
endpts[i,2] <- ybar + qnorm(.975,0,1)*S/sqrt(50)
if ((endpts[i,1] <= 10) && (10 <= endpts[i,2]))
  nin <- nin + 1
}
plot(c(endpts[1:100,1],endpts[1:100,2]),c(1:100,1:100),
     xlab="confidence intervals",ylab="number of trial")
for (i in 1:100) lines(c(endpts[i,1],endpts[i,2]),c(i,i))
lines(c(10,10),c(1,100))
return(nin)
}

```

Enter the code line by line. When you get no error messages, you call the function with an input command:

```
confdemo()
```

This will generate the graphical display and the output to the R console.

Lab Questions

1. Look over the R code carefully and try to figure out what each line does. Explain. In particular, how are the plotting commands (`plot` and `lines`) working? Note: From the first lab last semester, recall `c` is the R constructor for a list. `c(a,b)` denotes a list with entries a, b .
2. Use the `confdemo` function to generate the plot showing one hundred 95% confidence intervals for the mean, generated from samples of size $n = 50$ from a normal population with $\mu = 10$ and $\sigma = 2$.
3. Explain what your plot shows. In particular, how many of the intervals contain $\mu = 10$ and how many do not? How does that relate to the 95% confidence level?
4. Will the plot *always* show exactly 95 intervals containing the population mean $\mu = 10$? Try calling the function several more times if you got exactly 95 the first time. Does it make sense to say 95% of the confidence intervals produced here contain $\mu = 10$? In what sense?
5. An R function can also take *input values* that are used when the function is called. The names go in the set of parentheses in the first line, and then you supply values for them when you call the function. Modify the `confdemo` function so that you can specify the **sample size** n , and the confidence level $1 - \alpha$. **Hints:** The first line of your new function definition can look like

```
confdemo2 <- function(sampsize,conflevel)
```

and every where you used the sample size $n = 50$ before, now you will put in the name of the input `sampsize` instead. Similarly everywhere that used the $\alpha/2 = .025$ from before, now you will want `confllevel/2` (this one is a bit trickier – you will need to recall what the `qnorm` function does). If you like, you can really jazz this up so that the μ and σ can be specified as well.

6. Suppose you now use `confdemo` to study the 95% confidence intervals for the mean, generated from samples of the new size $n = 70$ from the same normal population with $\mu = 10$ and $\sigma = 2$. What will change (apart from the fact that the computation will take somewhat longer)? When you check your intuition, be sure to look carefully at the horizontal axis scales.
7. Suppose you now use `confdemo` to study the 99% confidence intervals for the mean, generated from samples of size $n = 50$ from the same normal population with $\mu = 10$ and $\sigma = 2$. What will change? When you check your intuition, be sure to look carefully at the horizontal axis scales.
8. A common *misconception* about confidence intervals can be stated as follows: “When you compute the 95% confidence interval for the mean from a particular sample, there’s a 95% chance that the population mean is contained in your interval.” This statement (taken literally) is actually *meaningless* – why? What could you do to modify the statement so that it makes sense and is a true statement about confidence intervals?
9. Another common *misconception* about confidence intervals can be stated as follows: “When you increase the sample size n , the width of the $(1 - \alpha) \times 100\%$ confidence interval always decreases.” This statement is actually false – why? What could you do to modify the statement so that it makes sense and is a true statement about confidence intervals?
10. A final common (and tempting) *misconception* about confidence intervals deals with the following situation. Say we are using the interval to decide whether evidence from samples supports the hypothesis that a population mean μ has a particular value μ_0 . If the value μ_0 is contained in a confidence interval but is far from the midpoint (even very close to one endpoint), it is tempting to think that the evidence indicates possibly $\mu \neq \mu_0$. This is an *incorrect* deduction in this situation. Look carefully back at your results. Explain why concluding $\mu \neq \mu_0$ on the basis of one confidence interval where μ_0 is close to an endpoint is *not a correct conclusion*.

Assignment

Lab reports containing input, output, and answers to the question posed above are due no later than Monday, February 27. If you do not finish during the hour today, you can return to HA 136 any time it is not in use by another class.