A Classic Bootstrap Example

Christopher A. T. Ferro

March 31, 2003

Consult the references given in Wilks (1997) to understand why the bootstrap works. Try the following, simple experiment to gain some confidence in the efficacy of bootstrap methods.

Generate from some distribution a sample of independent random numbers, $x_1, \ldots, x_n$, with mean

$$\bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i.$$

The usual estimate of the standard error of the mean, as used by the one-sample t-test for example, is $s/\sqrt{n}$, where

$$s^2 = \frac{1}{n-1} \sum_{i=1}^{n} (x_i - \bar{x})^2 .$$

Another estimate of the standard error can be obtained using bootstrap resampling. First choose the number of bootstrap samples, say $B = 1000$, noting that estimates will be more precise with larger B. Then implement the following algorithm.

do $b = 1, \ldots, B$
    do $i = 1, \ldots, n$
        choose x_i^* at random from $x_1, \ldots, x_n$
    end do
    calculate $\bar{x}_b^* = n^{-1} \sum_{i=1}^{n} x_i^*$
end do

The standard error of the mean may be estimated by $\sqrt{v_B}$, where

$$v_B = \frac{1}{B-1} \sum_{b=1}^{B} (\bar{x}_b^* - \bar{x}^*)^2 \qquad \text{and} \qquad \bar{x}^* = \frac{1}{B} \sum_{b=1}^{B} \bar{x}_b^*.$$

Compare v_B to s^2/n. You will find that they are similar if the sample size, n, is not too small. In fact, it is possible to show mathematically that v_B approximates the estimate $\tilde{s}^2/n$, where

$$\tilde{s}^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2 \, ,$$

with the approximation getting better, on average, for larger B.

Both v_B and s^2/n are estimates: neither is the true value. Both will get increasingly close to the truth as n and, for the bootstrap estimate, B increase. Similarly, both will perform poorly if the sample size, n, is small. The bootstrap approach is simply an alternative to standard methods. The advantage of the bootstrap is that it extends easily to other situations for which standard methods may not exist; on the other hand, it is not guaranteed to perform well in all situations, so some care is still required.

R Code

The following R code illustrates the example for random numbers generated from a uniform distribution on the interval $[0, 1]$. You can run the code on Rweb: follow links from "www.sciviews.org/_rgui/".

```
n <- 30          # sample size
B <- 1000        # number of bootstrap samples
x <- runif(n)    # random sample, size n, uniform on [0,1]
m <- mean(x)     # sample mean
mb <- rep(NA,B) # memory allocation

for(b in 1:B) {
  xb <- sample(x,n,replace=T) # resample with replacement
  mb[b] <- mean(xb)           # mean of the bootstrap sample
}

sqrt(var(x)/n) # estimated standard error
sqrt(var(mb))  # bootstrap estimated standard error
sqrt(1/(12*n)) # true standard error
```

References

Wilks, D.S. (1997) Resampling hypothesis tests for autocorrelated fields. *J. Climate* **10**, 65–82.