Bootstrap comparison of means between two groups

Crispin Jordan

04/09/2021

Overview

Johnston & Faulkner (2021)'s paper, "A bootstrap approach is a superior statistical method for the comparison of non-normal data with differing variances" (New Phytologist; https://doi.org/10.1111/nph.17159) provides a method to compare median values between two groups. As the title implies, their method may be suitable when data are not normally distributed and variances differ between the groups.

This document provides an example analysis using their method. We suggest you read the original publication for complete details to determine whether this method suits your purposes. For example, the article indicates that the method is robust for sample sizes of 10 or greater per group.

Code from Johnston & Faulkner (2021)

Johnston & Faulkner (2021) provide **R** code to implement their method. Their code includes:

```r
#### A bootstrap approach is a superior statistical method for the comparison of
#### non-normal data with differing variances
#### MG Johnston, C Faulkner
#### 15 December 2020

mcp_ci <- function(success, trials, alpha){
  ## Copyright (C) 2001 Frank E Harrell Jr
  ## Modified by Matthew G Johnston 2020 from binconf in the Hmisc package
  ## Distributed under GNU General Public License v2 (or later) without any
  ## warranty. See the GNU General Public License for more details.
  zcrit <-  - qnorm(alpha/2)
  z2 <- zcrit * zcrit
  mc_p <- success/trials
  cl <- (mc_p + z2/2/trials + c(-1, 1) * zcrit *
           sqrt((mc_p * (1 - mc_p) + z2/4/trials)/trials))/(1 + z2/trials)
  if(success == 1)
    cl[1] <-  - log(1 - alpha)/trials
  if(success == (trials - 1))
    cl[2] <- 1 + log(1 - alpha)/trials
  return(cbind(mc_p, lower_ci=cl[1], upper_ci=cl[2]))
}

medianBootstrap<- function(data1, data2, N=5000, alpha=0.05){
  ## Calculate observed test statistic
  mediandiff<-median(data1)-median(data2)
  ## Generate the null distribution
  boots<-replicate(N, median(sample(data1,length(data1), replace=T))
                 -median(sample(data2,length(data2),  replace=T))-mediandiff)
  ## Count the number of at resampled observations which are at least as extreme
```

```
  above <- sum(abs(boots)>=abs(mediandiff))
  ## Calculate p value and confidence intervals
  mcp<-mcp_ci(above+1,N+1, alpha)
  return(mcp)
}
```

An example analysis

Let's analyze the data we explored in the t-test Practice Problems. We begin by obtaining the data:

```
example.df <- read.table("MS1_Data.csv", header = TRUE, sep = ',')
```

Let's look at top of the dataframe:

```
head(example.df)
```

```
##   Sex Adj..Height Adj..Weight      BMI Waist.meas Hip.meas Waist.Hip.Ratio.WHpR
## 1   F       1.565          59 24.08925         73       97            0.7525773
## 2   F       1.615          54 20.70374         63       84            0.7500000
## 3   F       1.645          61 22.54229         80       88            0.9090909
## 4   F       1.745          56 18.39065         71       87            0.8160920
## 5   F       1.675          48 17.10849         64       84            0.7619048
## 6   M       1.855          77 22.37705         84       89            0.9438202
##   Waist.Height.Ratio.WHtR Dom.Hand Dom.Grip Non.dom.Grip dom.non.dom.ratio
## 1               0.4591195        R     27.0         25.0         1.0800000
## 2               0.3841463        R     21.0         21.0         1.0000000
## 3               0.4790419        R     27.0         26.0         1.0384615
## 4               0.4011299        R     22.0         20.0         1.1000000
## 5               0.3764706        R     20.0         17.0         1.1764706
## 6               0.4468085        R     53.2         53.6         0.9925373
##   Peak.Exp.Flow Sys.BP Dias.BP Heart.Rate
## 1           450    109      78         86
## 2           350    109      75         80
## 3           300    120      65         86
## 4           440     99      66         75
## 5           420     88      61         73
## 6           550    131      69         86
```

As we did in the t-test Practice Problems, we will compare `Peak.Exp.Flow` between females and males (`Sex`).

Let's first plot the data:

```
boxplot(Peak.Exp.Flow ~ Sex, data = example.df)
stripchart(Peak.Exp.Flow ~ Sex, data = example.df, vertical = TRUE, method =
             "jitter", pch = 21, col = "maroon", bg = "bisque", add = TRUE)
```

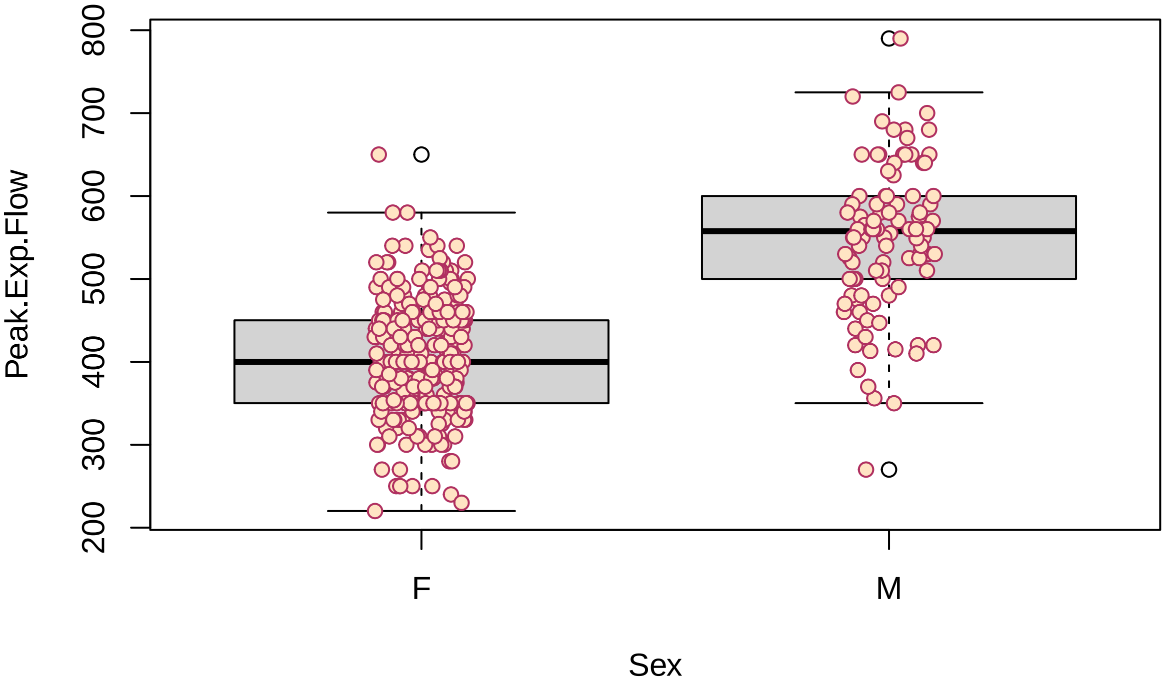

As described in the answers to the t-test Practice Problems, these data are suitable for analysis by a Welch's t-test.

Let's use the bootstrap approach to determine evidence for differences between medians of the two groups. First, we must create two vectors that contain the data from the two groups. We can do so like this:

```
Peak.Exp.Flow.female <- example.df[which(example.df$Sex == "F"),"Peak.Exp.Flow"]
Peak.Exp.Flow.male <- example.df[which(example.df$Sex == "M"),"Peak.Exp.Flow"]
```

To perform the analysis, we use the `medianBootstrap()` function from above. We provide the names of the two vectors that contain the data, we indicate the number of simulated datasets, and the type 1 error rate:

```
medianBootstrap(Peak.Exp.Flow.female, Peak.Exp.Flow.male, N=5000,
                           alpha=0.05)
```

```
##             mc_p     lower_ci    upper_ci
## [1,] 0.00019996 1.025661e-05 0.001131863
```

The output provides both a p-value (`mc_p`), as well as 95% confidence intervals for the p-value.

The output does not provide estimates of effect size with 95% confidence intervals. We will consider computational methods to obtain such estimates in the future.