Practical Stats Newsletter for April 2017

Subscribe and Unsubscribe: http://practicalstats.com/news
Archive of past newsletters http://www.practicalstats.com/news/bydate.html

In this newsletter:
A. Practical Stats Courses
B. How do Permutation Tests and Bootstrapping differ?
C. Coming Changes to Our Website

A. Practical Stats Courses
Our Applied Environmental Statistics courses are now on our online training site:
http://practicalstats.teachable.com/
The two courses separately are each $650 for a 1-year access for one person. Or get both courses together in a bundle for $1200. Note that the AES2 course (regression and trends) does not have all of its videos on the training site yet, so if you're planning to take only that course, don't sign up for it quite yet. The AES1 course is set, and if you get the bundle and start with 1, the videos for 2 will be there before you're ready for them.

The online short course "Use Of Practical Statistics in Site Investigations and Risk Assessments" will be broadcast live on May 16 and 18, 2017. Registration and information is through GeoEnviroPro. Cost: $138 CAD. See http://practicalstats.com/training/ for details.

B. How do Permutation Tests and Bootstrapping differ?
Permutation tests and bootstrapping are two procedures that are featured in our Applied Environmental Statistics online courses, and will get a future online course of their own this summer. What do each of these accomplish, and how do they differ?

Bootstrapping: Its purpose is to describe the variability of an estimate such as the mean. For that case the end result is a confidence interval. This can of course be done with t-intervals, but these assume that data follow a normal distribution (or that there is a sufficiently large sample size – see our July 2015 newsletter for more detail). A bootstrapped interval is computed without assuming a theoretical distributional shape for the data, allowing intervals to be computed for strongly-skewed data with outliers, both common characteristics of environmental data. Assuming a normal distribution is generally unrealistic for environmental data, due to asymmetry and outliers. Bootstrapped intervals avoid the pitfalls of assuming something that isn't true, providing more accurate interval estimates.

To compute a bootstrap interval, the data are re-sampled thousands of times. The process is called resampling with replacement, which is like copying a cell of a worksheet to another cell. The original cell remains intact and available for copying again. If the original dataset has 25 observations, each observation has a 1/25 likelihood of being selected at every step of the resampling. The first resample is to fill in spaces for 25

values, all of which come from the original observations, but not necessarily in the same proportions. One observation might randomly be chosen twice and copied twice to the new set of 25. Another observation may not be chosen at all. After the first resample of 25 is constructed, its mean is computed and saved. Then the process is repeated another 9998 or similarly large number of times. At the end, there are (in this case) 10,000 estimates of the mean (including the original one from the observed data). All 9999 of the means from the resamples are considered just as likely to have occurred as the mean that was actually observed. Taken together and ordered from smallest to largest, the mean at the position 10,000*0.025 = 250th from the bottom end is the low point on a two- sided 95% confidence interval -- 2.5% of the estimates are at or below that value. Similarly, the upper end is the mean at the position 10,000*0.975 = 975th from the bottom end of the ordered means. There are 2.5% of the estimates at or above that value. These comprise the two-sided 95% bootstrap confidence interval on the mean, and the interval will exhibit skewness if the original data were skewed. A t-interval is by definition symmetric, resulting in many cases with a lower limit that is too low, and lower than the bootstrap interval's limit.

Bootstrapping is used for computing confidence intervals around a regression equation, computing tolerance intervals, and many other purposes when defining the variability of a data characteristic or model.

Permutation Tests: Their purpose is to use resampling to avoid an assumption of a distributional shape when conducting hypothesis tests. So rather than the bootstrap objective of estimating intervals, permutation tests perform hypothesis tests. They differ from bootstrapping in that the resampling used is without replacement. In other words, all of 25 (say) observations are copied and their sequence is rearranged in a new column. The rearrangement breaks up the association between the observation value and the group it is assigned to. The collection of many rearrangements, with group assignments randomized, is collectively a picture of the test's null hypothesis. Permutation tests assume that any observed value could have come from any of the groups if the null hypothesis is true. This is called "exchangeability", another way of stating the null hypothesis -- all groups have the same distributional shape, so that all observations have come from that one distribution. Several past newsletters have discussed the usefulness of permutation tests as a method for testing differences in group means without assuming a normal distribution.

In addition to permutation tests as alternatives to group tests such as the t-test and ANOVA, permutation tests are used to test for serial correlation and other tests of "structure versus random occurrence". Nonparametric tests are special cases of permutation tests. Both methods are part of the class of procedures called "Resampling statistics", to which also belong 'jackknife' or 'leave-one-out' methods, the most commonly used of which is probably the PRESS statistic in regression.

C. Coming Changes to Our Website

I've never understood why businesses announce that their website will be changing. Just do it. I think the announcement is a veiled statement that something you wanted to be

completed by now is not yet finished. But here goes -- ours is being simplified, and will become "responsive", so able to be seen well from a tablet or phone. So look for a "fresh, new" looking website, which hopefully will also be easier to use, soon.

'Til next time,

Practical Stats -- Make sense of your data