Program name: artoutlierdetector

Usage: artoutlierdetector [-h -v -z -p *p-value*] -o *output prefix* -m *mask image* *input images ...*

-h or -help: Prints help message.

-v or -verbose: Enables verbose mode.

-z or -zero: Expands the mask to exclude from analysis any voxel that has the value of zero in any of the input images. That is, if a given image has a value of zero at a given voxel, that voxel will be excluded from all subsequent analyses in all images. Essentially, this option instructs the program to exclude voxels with missing (zero value) data.

-p or -pval: Specifies the p-value for outlier detection. Default is $p = 0.05$.

-o *output prefix*: Specifies a prefix used to name the output files.

-m or -mask *mask image*: Specifies a *mask image* of type *short* in ANALYZE format. Voxels in this image with values less than or equal to 0 are masked (excluded) from subsequent analyses. That is, the positive voxels in *mask image* are the voxels on which analysis takes place. All other voxels are simply ignored. This is a required argument.

input images ... are 2 or more images of type *short* in ANALYZE format. The recommended minimum number of input images is 10. The program attempts to detect and flag any outlier images in the input image set using the Grubbs' test for outlier detection (described below).

Program output: This program outputs 3 images: *prefix*_mask, *prefix*_avg, and *prefix*_sd, corresponding to the expanded mask (if the -z option is present), and the average and the standard deviation of the input images, respectively. A file entitled *prefix*_output.txt will also be created containing the program's text output, which amongst other things, indicates any outlier images. In addition to detecting the outliers, the program ranks all input images according to their closeness to the average image based on the y statistic defined below. Furthermore, the program outputs a ranking of the images according to their number of zero-valued voxels (missing data) that are not part of the specified mask. This output is only produced with the -z option.

Theory: The computer program *artoutlierdetector* uses the Grubbs' test for outliers defined for the hypothesis (Grubbs, 1969):

H_0 : There are no outliers in the input image set

H_1 : There is at least one outlier in the image set

If H_0 is rejected, the image with the largest Grubbs' test statistic G is declared an outlier. This image is then removed from the data set and the test procedure is repeated on the remaining images until no more outliers are detected.

Let n be the number of subjects (i.e., the number of images specified as input to the program), and p be the number of non-masked voxels in each image. Thus, the data can be summarized in $n \times p$ matrix $\mathbf{X}$. Each row of this matrix represents the non-masked voxels from a given subject. Thus, an element x_{ij} of matrix $\mathbf{X}$ corresponds to the jth voxel value of the ith subject.

The *artoutlierdetector* program first computes the data mean and variance at each voxel j as:

$$x_{\bullet j} = \frac{1}{n} \sum_{i=1}^{n} x_{ij} \tag{1}$$

and

$$\sigma_j^2 = \frac{1}{n-1} \sum_{i=1}^{n} (x_{ij} - x_{\bullet j})^2. \tag{2}$$

It will then compute a χ^2-like statistic (y_i) for each subject i as follows:

$$y_i = \sum_{j=1}^{p} \frac{(x_{ij} - x_{\bullet j})^2}{\sigma_j^2} \tag{3}$$

For large p, the y_i are approximately normally distributed. The sample mean ($\bar{y}$) and standard deviation (s) of y_i are computed as: $\bar{y} = (1/n) \sum_{i=1}^{n} y_i$, and

$$s = \left[\frac{1}{n-1} \sum_{i=1}^{n} (y_i - \bar{y})^2 \right]^{1/2}.$$

The Grubbs' test statistic G is then defined as:

$$G = \frac{y_{max} - \bar{y}}{s},$$

where

$$y_{max} = \max_i y_i.$$

The one-sided null hypothesis H_0 is rejected if:

$$G > \frac{n-1}{\sqrt{n}} \sqrt{\frac{t^2_{(\alpha/n,n-2)}}{n-2+t^2_{(\alpha/n,n-2)}}},$$

where $t_{(\alpha/n,n-2)}$ denotes the α/n critical value of the Student's t-distribution with $(n-2)$ degrees of freedom, and α is the p-value specified by the user (using the -p option; default=0.05). Note that this program only performs a one-sided test, because we are only interested in images with large y_i, i.e., those that are 'far' from the mean image as determined by Eq. 3.

References:

Grubbs, F. (1969). Procedures for Detecting Outlying Observations in Samples. *Technometrics*, Vol. 11, No. 1, pp. 1–21.

Program and documentation author: Babak A. Ardekani, Ph.D., Center for Advanced Brain Imaging, The Nathan S. Kline Institute for Psychiatric Research.

Oct. 9, 2007