

PAST

PAleontological SStatistics

Version 1.93



Reference manual

Øyvind Hammer

Natural History Museum

University of Oslo

ohammer@nhm.uio.no

1999-2009

Welcome to the PAST!

This program was originally designed as a follow-up to PALSTAT, a software package for paleontological data analysis written by P.D. Ryan, D.A.T. Harper and J.S. Whalley (Ryan et al. 1995).

Through continuous development for more than ten years, PAST has grown into a comprehensive statistics package that is used not only by paleontologists, but in many fields of life science, earth science, and even engineering and economics.

Further explanations of many of the techniques implemented together with case histories are found in Harper (1999). Also, "Paleontological data analysis" (Hammer & Harper 2005) can be viewed as a companion book to PAST.

If you have questions, bug reports, suggestions for improvements or other comments, we would be happy to hear from you. Contact us at ohammer@nhm.uio.no. For bug reports, remember to send us the data used, as saved from PAST, together with a complete description of the actions that lead to the problem.

The latest version of Past, together with documentation and a link to the Past mailing list, are found at

<http://folk.uio.no/ohammer/past>

We are grateful if you cite PAST in scientific publications. The official reference is Hammer et al. (2001).

References

Hammer, Ø. & Harper, D.A.T. 2006. Paleontological Data Analysis. Blackwell.

Hammer, Ø., Harper, D.A.T., and P. D. Ryan, 2001. PAST: Paleontological Statistics Software Package for Education and Data Analysis. *Palaeontologia Electronica* 4(1): 9pp.

Harper, D.A.T. (ed.). 1999. Numerical Palaeobiology. John Wiley & Sons.

Installation

The installation of PAST is easy: Just download the file 'Past.exe' and put it anywhere on your harddisk. Double-clicking the file will start the program. Windows will consider this a breach of security, and will ask if you trust the software provider. If you want to use the program, you will have to answer yes.

We suggest you make a folder called 'past' anywhere on your harddisk, and put all the files in this folder.

Please note: Problems have been reported for non-standard default font size in Windows - it may be necessary for the user to increase the sizes of windows in order to see all the text and buttons. If this happens, please set the font size to 'Small fonts' in the Screen control panel in Windows.

When you exit PAST, a file called 'pastsetup' will be automatically placed in your personal folder (for example 'My Documents' in Windows 95/98), containing the last used file directories.

The lack of "formal" Windows installation is intentional, and allows installation without administrator privileges.

The spreadsheet and the Edit menu

PAST has a spreadsheet-like user interface. Data are entered as an array of cells, organized in rows (horizontally) and columns (vertically).

Entering data

To input data in a cell, click on the cell with the mouse and type in the data. This can only be done when the program is in the 'Edit mode'. To select edit mode, tick the box above the array. When edit mode is off, the array is locked and the data cannot be changed. The cells can also be navigated using the arrow keys.

Any text can be entered in the cells, but almost all functions will expect numbers. Both comma (,) and decimal point (.) are accepted as decimal separators.

Absence/presence data are coded as 0 or 1, respectively. Any other positive number will be interpreted as presence. Absence/presence-matrices can be shown with black squares for presences by ticking the 'Square mode' box above the array.

Genetic sequence data are coded using C, A, G, T and U (lowercase also accepted).

Missing data are coded with question marks ('?') or the value -1. Unless support for missing data is specifically stated in the documentation for a function, the function *will not handle missing data correctly*, so be careful.

The convention in PAST is that items occupy rows, and variables columns. Three brachiopod individuals might therefore occupy rows 1, 2 and 3, with their lengths and widths in columns A and B. Cluster analysis will always cluster *items*, that is rows. For Q-mode analysis of associations, samples (sites) should therefore be entered in rows, while taxa (species) are in columns. For switching between Q-mode and R-mode, rows and columns can easily be interchanged using the Transpose operation.

Selecting areas

Most operations in PAST are only carried out on the area of the array which you have *selected* (marked). If you try to run a function which expects data, and no area has been selected, you will get an error message.

- A row is selected by clicking on the row label (leftmost column).
- A column is selected by clicking on the column label (top row).
- Multiple rows are selected by selecting the first row label, then shift-clicking (clicking with the Shift key down) on the additional row labels. Note that you can not 'drag out' multiple rows - this will instead move the first row (see below).
- Multiple columns are similarly marked by shift-clicking the additional column labels.

- The whole array can be selected by clicking the upper left corner of the array (the empty grey cell) or by choosing 'Select all' in the Edit menu.
- Smaller areas within the array can be selected by 'dragging out' the area, but this only works when 'Edit mode' is off.

IMPORTANT: Unfortunately, you can not select several columns that are not neighbours. This means that if you want to run an analysis on say the first and the third columns only, you will first have to move the columns together (see next paragraph).

Moving a row or a column

A row or a column (including its label) can be moved simply by clicking on the label and dragging to the new position.

Renaming rows and columns

When PAST starts, rows are numbered from 1 to 99 and columns are labelled A to Z. For your own reference, and for proper labelling of graphs, you should give the rows and columns more descriptive but short names. Choose 'Rename columns' or 'Rename rows' in the Edit menu. You must select the whole array, or a smaller area as appropriate.

Another way is to select the 'Edit labels' option above the spreadsheet. The first row and column are now editable in the same way as the rest of the cells.

Increasing the size of the array

By default, PAST has 99 rows and 26 columns. If you should need more, you can add rows or columns by choosing 'Insert more rows' or 'Insert more columns' in the Edit menu. Rows/columns will be inserted after the marked area, or at the bottom/right if no area is selected. When loading large data files, rows and/or columns are added automatically as needed.

Cut, copy, paste

The cut, copy and paste functions are found in the Edit menu. You can cut/copy data from the PAST spreadsheet and paste into other programs, for example Word and Excel. Likewise, data from other programs can be pasted into PAST – these need to be in a tab-separated text format.

Remember that local blocks of data (not all rows or columns) can only be marked when 'Edit mode' is off.

All modules giving graphic output have a 'Copy graphic' button. This will place the graphical image into the paste buffer for pasting into other programs, such as a drawing program for editing the

image. Graphics are copied using the 'Enhanced Metafile Format' in Windows. This allows editing of individual image elements in other programs. When pasting into Coreldraw, you have to choose 'Paste special' in the Edit menu, and then choose 'Enhanced metafile'. Some programs may have idiosyncratic ways of interpreting EMF images - beware of funny things happening.

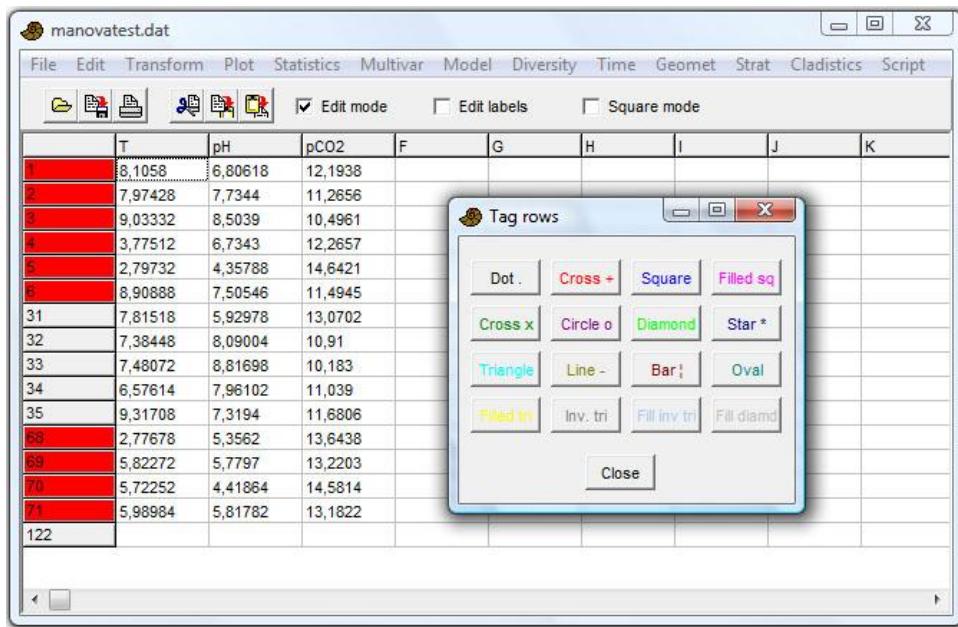
Remove

The remove function (Edit menu) allows you to remove selected row(s) or column(s) from the spreadsheet. The removed area is not copied to the paste buffer.

Grouping (coloring) rows

Selected rows (data points) can be tagged with one of 16 attractive colors using the 'Row color/symbol' option in the Edit menu. Each group is also associated with a symbol (dot, cross, square, diamond, plus, circle, triangle, line, bar, filled square, star, oval, filled triangle, inverted triangle, filled inverted triangle, filled diamond). This is useful for showing different groups of data in plots, and is also required by a number of analysis methods.

Important: For methods that require groupings of rows using colors, rows belonging to one group must be consecutive. If more than 16 groups are required, colors can be re-used. In the example below, three groups have been correctly marked.



The 'Numbers to colors' option in the Edit menu allows the numbers 1-16 in one selected column to set corresponding colors (symbols) for the rows.

Selecting datatypes for columns

Selected columns can be tagged with a datatype (continuous/unspecified, ordinal, nominal or binary) using the 'Column data types' option in the Edit menu. This is only required if you wish to use mixed similarity/distance measures.

Transpose

The Transpose function, in the Edit menu, will interchange rows and columns. This is used for switching between R mode and Q mode in cluster analysis, principal components analysis and seriation.

Grouped columns to multivar

Converts from a format with multivariate items presented in consecutive groups of N columns to the Past format with one item per row and all variates along the columns. For $N=2$, two specimens and four variables a-d, the conversion is from

a₁ b₁ a₂ b₂
c₁ d₁ c₂ d₂

to

a₁ b₁ c₁ d₁
a₂ b₂ c₂ d₂

Grouped rows to multivar

Converts from a format with multivariate items presented in consecutive groups of N rows to the Past format with one item per row and all variates along the columns. For $N=2$, two specimens and four variables a-d, the conversion is from

a₁ b₁
c₁ d₁
a₂ b₂
c₂ d₂

to

a₁ b₁ c₁ d₁
a₂ b₂ c₂ d₂

Stack colored rows into columns

Stacks colored groups horizontally along columns. This can be useful e.g. for performing univariate statistics on pairs of columns across groups.

Samples to events (UA to RASC)

Given a data matrix of occurrences of taxa in a number of samples in a number of sections, as used by the Unitary Associations module, this function will convert each section to a single row with orders of events (FADs, LADs or both) as expected by the Ranking-Scaling module. Tied events (in the same sample) will be given equal ranking.

Events to samples (RASC to UA)

Expects a data matrix with sections/wells in rows, and taxa in columns, with FAD and LAD values in alternating columns (i.e. two columns per taxon). Converts to the Unitary Associations presence/absence format with sections in groups of rows, samples in rows and taxa in columns.

Loading and saving data

The 'Open' function is found in the File menu. You can also drag a file from the desktop onto the PAST window. PAST uses an ASCII file format, for easy importing from other programs (e.g. Word) and easy editing in a text editor. The format is as follows:

```
.      columnlabel columnlabel columnlabel  
rowlabel data      Data      Data  
rowlabel data      Data      Data  
rowlabel data      Data      Data
```

Empty cells (like the top left cell) are coded with a full stop (.). Cells are separated by white space, which means that you *must never use spaces in row or column labels*. 'Oxford Clay' is thus an illegal column label which would confuse the program.

If any rows have been assigned a color other than black, the row labels in the file will start with an underscore, a number from 0 to 11 identifying the color (symbol), and another underscore.

If any columns have been assigned a datatype other than continuous/unspecified, the column labels in the file will similarly start with an underscore, a number from 0-3 identifying the datatype (0=continuous/unspecified, 1=ordinal, 2=nominal, 3=binary), and another underscore.

In addition to this format, Past can also detect and open files in the following formats:

- Excel (only the first worksheet).

- Nexus format (see below), popular in systematics.
- TPS format developed by Rohlf (the landmark, outlines, curves, id, scale and comment fields are supported, other fields are ignored).
- BioGraph format for biostratigraphy (SAMPLES or DATUM format). If a second file with the same name but extension '.dct' is found, it will be included as a BioGraph dictionary.
- RASC format for biostratigraphy. You must open the .DAT file, and the program expects corresponding .DIC and .DEP files in the same directory.
- CONOP format for biostratigraphy. You must open the .DAT file (log file), and the program expects corresponding .EVT (event) and .SCT (section) files in the same directory.

The 'Insert from file' function is useful for concatenating data sets. The loaded file will be inserted into your existing spreadsheet at the selected position (upper left). Other data sets can thus be inserted both to the right of and below your existing data.

Data from Excel

Data from Excel can be imported in three ways:

- Copy from Excel and paste into PAST. Note that if you want the first row and column to be copied into the label cells in PAST, you must switch on the "Edit labels" option.
- Open the Excel file from PAST. The "Edit labels" option operates in the same way.
- Make sure that the top left cell in Excel contains a single dot (.) and save as tab-separated text in Excel. The resulting text file can be opened directly in PAST.

Reading and writing Nexus files

The Nexus file format is used by many systematics programs. PAST can read and write the Data (character matrix) block of the Nexus format. Interleaved data are supported. Also, if you have performed a parsimony analysis and the 'Parsimony analysis' window is open, all shortest trees will be written to the Nexus file for further processing in other programs (e.g. MacClade or Paup). Note that not all Nexus options are presently supported.

Counter

A counter function is available in the Edit menu for use e.g. at the microscope when counting microfossils of different taxa. A single row (sample) must be selected. The counter window will open with a number of counters, one for each selected column (taxon). The counters will be initialized with the column labels and any counts already present in the spreadsheet. When closing the counter window, the spreadsheet values will be updated.

Count up (+) or down (-) with the mouse, or up with the keys 1-9, 0 and a-z (only the first 36 counters). The bars represent relative abundance. A log of events is given at the far right - scroll up and down with mouse or arrow keys inside the text. An optional auditive feedback has a specific pitch for each counter.

Transform menu

These routines subject your data to mathematical operations. This can be useful for bringing out features in your data, or as a necessary preprocessing step for some types of analysis.

Logarithm

The Log function in the Transform menu log-transforms your data using the base-10 logarithm. If the data contain zero or negative values, it may be necessary to add a constant (e.g. 1) before log-transforming (use Evaluate Expression $x+1$).

This is useful, for example, to compare your sample to a log-normal distribution or for fitting to an exponential model. Also, abundance data with a few very dominant taxa may be log-transformed in order to downweight those taxa.

Missing data supported.

Remove trend

This function removes any linear trend from a data set (two columns with X-Y pairs, or one column with Y values). This is done by subtraction of a linear regression line from the Y values. Removing the trend can be a useful operation prior to time series analyses such as spectral analysis, auto- and cross-correlation and ARMA.

Subtract mean

This function subtracts the column mean from each of the selected columns. The means cannot be computed row-wise.

Missing values supported.

Row percentage

All values converted to the percentage of the row sum.

Missing values supported.

Row normalize length

All values divided by the Euclidean length of the row vector.

Missing values supported.

Abundance to presence/absence

All positive (non-zero) numbers are replaced with ones.

Missing values supported.

Procrustes fitting

Transforms your measured point coordinates to Procrustes coordinates. There is also a menu choice for Bookstein coordinates. Specimens go in different rows and landmarks along each row. If you have three specimens with four landmarks in 2D, your data should look as follows:

```
x1 y1 x2 y2 x3 y3 x4 y4  
x1 y1 x2 y2 x3 y3 x4 y4  
x1 y1 x2 y2 x3 y3 x4 y4
```

For 3D the data will be similar, but with additional columns for z.

Landmark data in this format could be analyzed directly with the multivariate methods in PAST, but it is recommended to standardize to Procrustes coordinates by removing position, size and rotation. A further transformation to Procrustes residuals (approximate tangent space coordinates) is achieved by selecting 'Subtract mean' in the Edit menu. You must convert to Procrustes coordinates first, then to Procrustes residuals.

The "Rotate to major axis" option places the result into a standard orientation for convenience.

A thorough description of Procrustes and tangent space coordinates is given by Dryden & Mardia (1998). The algorithms for Procrustes fitting are from Rohlf & Slice (1990) (2D) and Dryden & Mardia (1998) (3D). It should be noted that for 2D, the iterative algorithm of Rohlf & Slice (1990) often gives slightly different results from the direct algorithm of Dryden & Mardia (1998). Past uses the former in order to follow the "industry standard".

Missing data is supported by column average substitution.

References

Dryden, I.L. & K.V. Mardia 1998. Statistical Shape Analysis. Wiley.

Rohlf, F.J. & Slice, D. 1990.

Bookstein fitting

Bookstein fitting has a similar function as Procrustes fitting, but simply standardizes size, rotation and scale by forcing the two first landmarks onto the coordinates (0,0) and (1,0). It is not in common use today. Bookstein fitting is only implemented for 2D.

Remove size from landmarks

The 'Remove size from landmarks' option in the Transform menu allows you to remove size by dividing all coordinate values by the centroid size for each specimen (Procrustes coordinates are also normalized with respect to size).

See Dryden & Mardia (1998), p. 23-26.

References

Dryden, I.L. & K.V. Mardia 1998. Statistical Shape Analysis. Wiley.

Transform landmarks

Allows rotation of the point cloud in steps of 90 degrees, and top-bottom or left-right flipping (mirroring), mainly for plotting convenience. The mirror operation may be useful for reducing a bilaterally symmetric landmark data, by Procrustes fitting the left half to a mirrored version of the right half (and optionally averaging the two).

Only for 2D coordinates.

Remove size from distances

Attempts to remove the size component from a multivariate data set of measured distances (specimens in rows, variables in columns). Three methods are available.

- *Isometric Burnaby's method* projects the set of measured distances onto a space orthogonal to the first principal component. Burnaby's method may (or may not!) remove isometric size from the data, for further "size-free" data analysis. Note that the implementation in PAST does not center the data within groups - it assumes that all specimens (rows) belong to one group.

- *Allometric Burnaby's method* will log-transform the data prior to projection, thus conceivably removing also allometric size-dependent shape variation from the data.
- *Allometric vs. standard* estimates allometric coefficients with respect to a standard (reference) measurement L such as overall length (Elliott et al. 1995). This standard variable should be placed in the first column. Each additional column is regressed onto the first column after log-transformation, giving a slope (allometric coefficient) b for that variable. An adjusted measurement is then computed from the original value M as

$$M_{adj} = M \left(\frac{\bar{L}}{L} \right)^b.$$

Reference

Elliott, N.G., K. Haskard & J.A. Koslow 1995. Morphometric analysis of orange roughy (*Hoplostethus atlanticus*) off the continental slope of southern Australia. *Journal of Fish Biology* 46:202-220.

Sort ascending and descending

Sorts the rows in the marked area, based on values in the selected data column.

The "Sort descending" function is useful, for example, to plot taxon abundances against their ranks (this can also be done with the Abundance Model module).

Sort on color

Sorts the rows in the marked area on color.

Column difference

Simply subtracts two selected columns, and places the result in the next column.

Regular interpolation

Moves an irregularly sampled time series or transect (possibly multivariate) onto a regular spacing, as required by most methods for time series analysis. The x values should be in the first selected column. These will be replaced by a regularly increasing series (with the original number of points). All additional selected columns will be interpolated linearly to match.

The perils of linear interpolation should be kept in mind.

Evaluate expression

This powerful feature allows flexible mathematical operations on the selected array of data. Each selected cell is evaluated, and the result replaces the previous contents. A mathematical expression must be entered, which can include any of the operators $+$, $-$, $*$, $/$, $^$ (power), and \bmod (modulo). Also supported are brackets $()$, and the functions abs , atan , cos , sin , exp , \ln , sqrt , sqr , round and trunc .

The following values are also defined:

- x (the contents of the current cell)
- l (the cell to the left if it exists, otherwise 0)
- r (the cell to the right)
- u (the cell above, or up)
- d (the cell below, or down)
- mean (the mean value of the current column)
- min (the minimum value)
- max (the maximum value)
- n (the number of cells in the column)
- i (the row index)
- j (the column index)
- random (uniform random number from 0 to 1)
- normal (Gaussian random number with mean 0 and variance 1).
- integral (running sum of the current column)
- stdev (standard deviation of the current column)
- sum (total sum of the current column)

In addition, other columns can be referred to using the column name preceded by ' $c_{}$ ', for example c_A .

Examples:

$\text{sqrt}(x)$	Replaces all numbers with their square roots
$(x-\text{mean})/\text{stdev}$	Mean and standard deviation normalization, column-wise
$x-0.5*(\text{max}+\text{min})$	Centers the values around zero
$(u+x+d)/3$	Three-point moving average smoothing
$x-u$	First-order difference
i	Fills the column with the row numbers (requires non-empty cells, such as all zeros)
$\text{sin}(2*3.14159*i/n)$	Generates one period of a sine function down a column (requires non-empty cells)
$5*\text{normal}+10$	Random number from a normal distribution, with mean of 10 and standard deviation of 5.

Missing values supported.

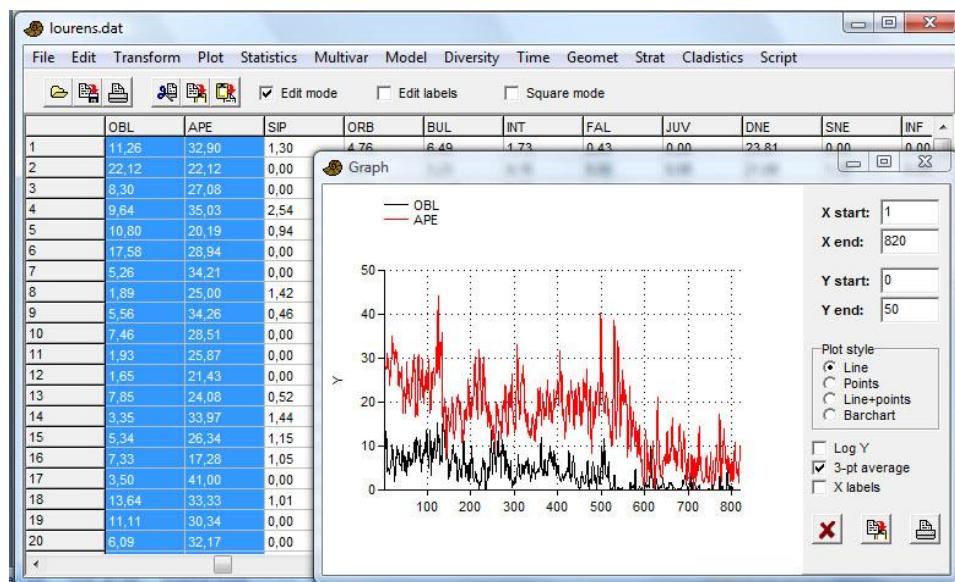
Plot menu

Graph

Plots one or more columns as separate graphs. The x coordinates are set automatically to 1,2,3,... There are four plot styles available: Graph (line), points, line with points, and bars. The 'X labels' option sets the x axis labels to the appropriate row names.

The "Log Y" option log-transforms the Y values only, to base 10, but with log 0 set to 0.

Missing values are disregarded.



XY graph

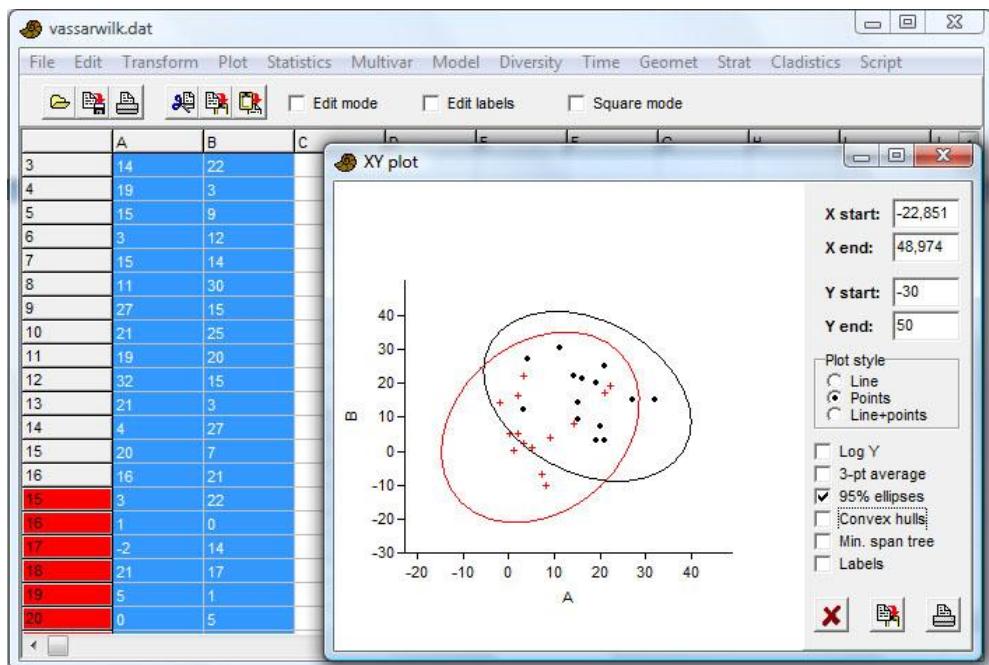
Plots one or more pairs of columns containing x/y coordinate pairs. The 'log Y' option log-transforms your Y values (if necessary, a constant is added to make the minimum log value equal to 0). The curve can also be smoothed using 3-point moving average.

95% concentration ellipses can be plotted in most scatter plots in PAST, such as scores for PCA, CA, DCA, PCO and NMDS. The calculation of these ellipses assumes bivariate normal distribution.

Convex hulls can also be drawn in the scatter plots, in order to show the areas occupied by points of different 'colors'. The convex hull is the smallest convex polygon containing all points.

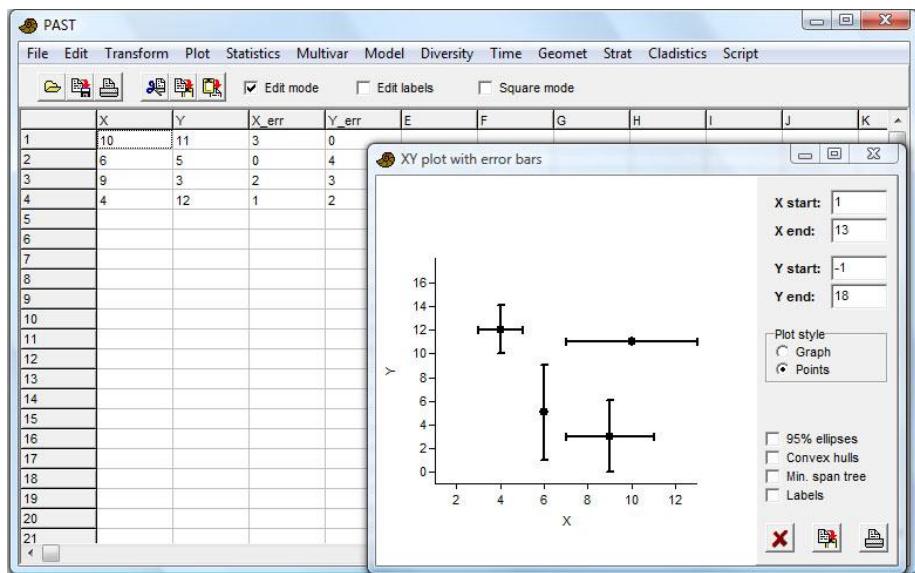
The minimal spanning tree is the set of lines with minimal total length, connecting all points. In the XY graph module, Euclidean lengths in 2D are used.

Hold the mouse cursor over a point to see its row label.



XY graph with error bars

As XY graph, but expects four columns (or a multiple), with x, y, x error and y error values. Symmetric error bars are drawn around each point. If an error value is set to zero, the corresponding error bar is not drawn.



Histogram

Plots histograms (frequency distributions) for one or more columns. The number of bins is by default set to an "optimal" number (the zero-stage rule of Wand 1997):

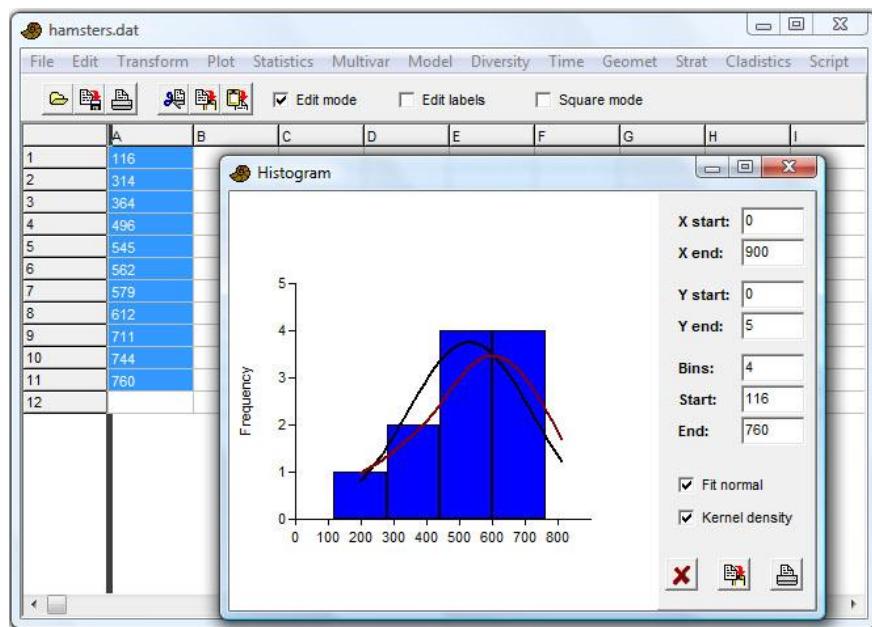
$$h = 3.49 \min(s, IQ/1.349) n^{-1/3}$$

where s is the sample standard deviation and IQ the interquartile range.

The number of bins can be changed by the user. The "Fit normal" option draws a graph with a fitted normal distribution (Parametric estimation, not Least Squares).

Kernel Density Estimation is a smooth estimator of the histogram. PAST uses a Gaussian kernel with range according to the rule given by Silverman (1986):

$$h = 0.9 \min(s, IQ/1.34) n^{-1/5}.$$



References

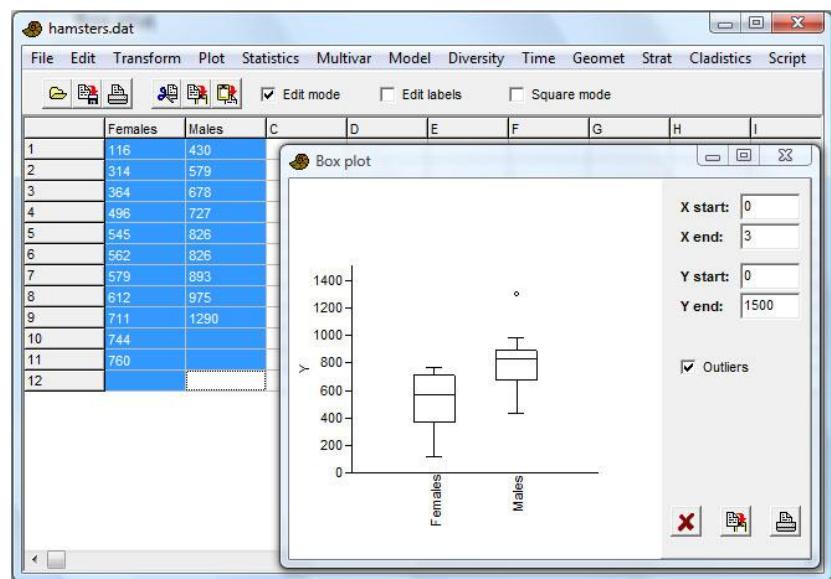
Silverman, B.W. 1986. Density estimation for statistics and data analysis. Chapman & Hall.

Wand, M.P. 1997. Data-based choice of histogram bin width. American Statistician 51:59-64.

Box plot

Box plot for one or several columns (samples) of univariate data. For each sample, the 25-75 percent quartiles are drawn using a box. The median is shown with a horizontal line inside the box. The minimal and maximal values are shown with short horizontal lines ("whiskers").

If the "Outliers" box is ticked, another box plot convention is used. The whiskers are drawn from the top of the box up to the largest data point less than 1.5 times the box height from the box (the "upper inner fence"), and similarly below the box. Values outside the inner fences are shown as circles, values further than 3 times the box height from the box (the "outer fences") are shown as stars.



Percentiles

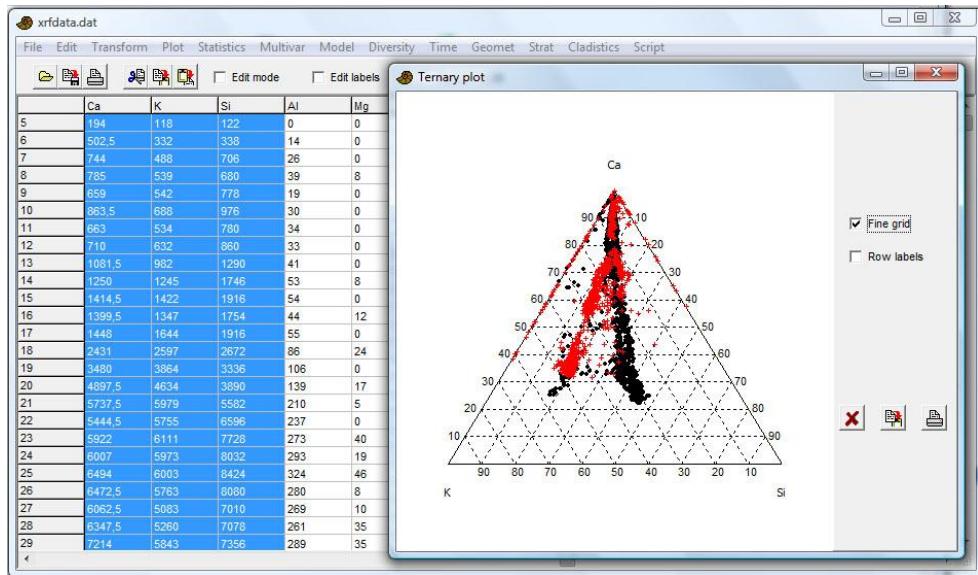
Plots the (cumulative) number of points smaller than each percentage of points. Two popular methods are included. For a percentile p , the rank is computed according to $k=p(n+1)/100$, and the value that corresponds to that rank taken. In the rounding method, k is rounded to the nearest integer, while in the interpolation method, non-integer ranks are handled by interpolation between the two nearest ranks.

Normal probability plot

Plots a normal probability (normal QQ) plot for one column of data. A normal distribution will plot on a straight line. For comparison, an RMA regression line is given, together with the Probability Plot Correlation Coefficient.

Ternary

Ternary plot for three columns of data, normally containing proportions of compositions.



Bubble plot

An attempt at plotting 3D data (three columns) by showing the third axis as size of disks. Negative values are not shown. Select "Subtract min" to subtract the smallest third axis value from all values - this will force the data to be positive. The "Size" slider scales the bubbles relative to unit radius on the x axis scale.

Survivorship

Survivorship curves for one or more columns of data. The data will normally consist of age or size values. A survivorship plot shows the number of individuals which survived to different ages. Assuming exponential growth (highly questionable!), size should be log-transformed to age. This can be done either in the Transform menu, or directly in the Survivorship dialogue.

Landmarks

This function is very similar to the 'XY graph', the only difference being that all XY pairs on each row are plotted with the appropriate row color and symbol. It is well suited for plotting landmark data.

Landmarks 3D

Plotting of points in 3D (XYZ triples). Especially suited for 3D landmark data, but can also be used e.g. for PCA scatter plots along three principal components. The point cloud can be rotated around the x and the y axes (note: left-handed coordinate system). The 'Perspective' slider is normally not used. The 'Stems' option draws a line from each point down to a bottom plane, which can sometimes enhance 3D information. 'Lines' draws lines between consecutive landmarks within each separate specimen (row). 'Axes' shows the three coordinate axes with the centroid of the points as the origin.

Matrix

Two-dimensional plot of the data matrix, using a grayscale with white for lowest value, black for highest. Can be useful to get an overview over a large data matrix.

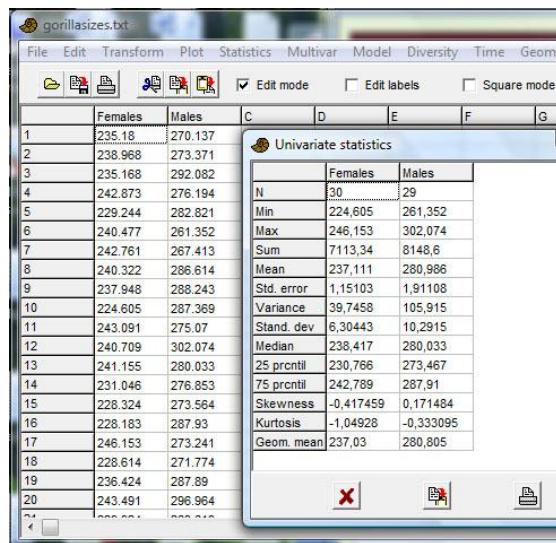
Surface

Three-dimensional landscape plot of a data matrix containing elevation values. Colors can be assigned according to height, or gray-shaded using a fixed light source.

Statistics menu

Univariate

This function computes a number of basic descriptive statistics for one or more samples of univariate data. Each sample must have at least 3 values, and occupies one column in the spreadsheet. The columns do not have to contain the same number of values. The example below uses two samples: the sizes in mm of the skulls of 30 female and 29 male gorillas. To run the analysis, the two columns (or the whole spreadsheet) must be selected.



The following numbers are shown for each sample:

N: The number of values n in the sample

Min: The minimum value

Max: The maximum value

Sum: The sum

$$\bar{x} = \frac{\sum x_i}{n}$$

Mean: The estimate of the mean, calculated as

$$SE_{\bar{x}} = \frac{s}{\sqrt{n}}$$
 where s is the estimate of the standard deviation (see below).

Std. error: The standard error of the estimate of the mean, calculated as

Variance: The sample variance, calculated as $s^2 = \frac{1}{n-1} \sum (x_i - \bar{x})^2$.

Stand. dev.: The sample standard deviation, calculated as $s = \sqrt{\frac{1}{n-1} \sum (x_i - \bar{x})^2}$.

Median: The median of the sample. For n odd, the given value such that there are equally many values above and below. For n even, the average of the two central values.

25 prctil: The 25th percentile, i.e. the given value such that 25% of the sample is below, 75% above. Linear interpolation is used when necessary.

75 prctil: The 75th percentile, i.e. the given value such that 75% of the sample is below, 25% above. Linear interpolation is used when necessary.

Skewness: The sample skewness, zero for a normal distribution, positive for a tail to the right.

$$G_1 = \frac{n}{(n-1)(n-2)} \frac{\sum (x_i - \bar{x})^3}{\left(\sqrt{\frac{1}{n-1} \sum (x_i - \bar{x})^2} \right)^3}$$

Calculated as

Note there are several versions of this around – Past uses the same equation as SPSS and Excel. Slightly different results may occur using other programs, especially for small sample sizes.

Kurtosis: The sample kurtosis, zero for a normal distribution. Calculated as

$$G_2 = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \frac{\sum (x_i - \bar{x})^4}{\left(\sqrt{\frac{1}{n-1} \sum (x_i - \bar{x})^2} \right)^4} - 3 \frac{(n-1)^2}{(n-2)(n-3)}$$

Again Past uses the same equation as SPSS and Excel.

Geom. mean: The geometric mean, calculated as $(x_1 x_2 \cdots x_n)^{1/n}$.

Missing data: Supported by deletion.

Similarity and distance indices

Computes a number of similarity or distance measures between all pairs of rows. The data can be univariate or (more commonly) multivariate, with variables in columns. The results are given as a symmetric similarity/distance matrix. This module is rarely used, because similarity/distance matrices are usually computed automatically from primary data in modules such as PCO, NMDS, cluster analysis and ANOSIM in Past.

	Pinna	Chlamm	Pleuro	Ctenos	Nanogy	Myopho	Chlams	Chlamq	Trautm	Discor
Loc1	0	0	61	0	767	0	0	0	0	0
Loc2	0	48	1304	0	656	164	0	0	0	0
Loc3	342	72	18	7	0	0	0	0	0	0
Loc4	0	0	12	0	127	0	0	0	0	0
Loc5	0	0	11	0	68	8	0	0	0	15
Loc6	0	0	0	0	985	316	106	46	42	36

	Loc1	Loc2	Loc3	Loc4	Loc5	Loc6
Loc1	1	0,478	0,028414	0,28749	0,16989	0,65028
Loc2	0,478	1	0,050555	0,12029	0,076517	0,44288
Loc3	0,028414	0,050555	1	0,041522	0,040665	0
Loc4	0,28749	0,12029	0,041522	1	0,6556	0,1521
Loc5	0,16989	0,076517	0,040665	0,6556	1	0,11145
Loc6	0,65028	0,44288	0	0,1521	0,11145	1

Index

- Gower
- Euclidean
- Correlation
- Rho
- Dice
- Jaccard
- Kulczynski
- Ochiai
- Simpson
- Bray-Curtis
- Cosine
- Morisita
- Raup-Crick
- Horn
- Hamming
- Chord
- Manhattan
- Jukes-Cantor
- Kimura
- Tajima-Nei

Gower

A distance measure that averages the difference over all variables, each term normalized for the range of that variable:

$$d_{jk} = \frac{1}{n} \sum_i \frac{|x_{ji} - x_{ki}|}{\max_s x_{si} - \min_s x_{si}}.$$

The Gower measure is similar to Manhattan distance (see below) but with range normalization. When using mixed data types (see below), this is the default measure for continuous and ordinal data.

Euclidean

Basic Euclidean distance. In early versions of Past, this was normalized for the number of variables (the value is still adjusted for missing data).

$$d_{jk} = \sqrt{\sum_i (x_{ji} - x_{ki})^2}$$

Correlation

The complement $1-r$ of Pearson's r correlation across the variables:

$$d_{jk} = 1 - \frac{\sum_i (x_{ji} - \bar{x}_j)(x_{ki} - \bar{x}_k)}{\sqrt{\sum_i (x_{ji} - \bar{x}_j)^2} \sqrt{\sum_i (x_{ki} - \bar{x}_k)^2}}$$

Taking the complement makes this a distance measure. See also the Correlation module, where Pearson's r is given directly and with significance tests.

Rho

The complement $1-r_s$ of Spearman's rho, which is the correlation coefficient of ranks. See also the Correlation module, where rho is given directly and with significance tests.

Dice

Also known as the Sorensen coefficient. For binary (absence-presence) data, coded as 0 or 1 (any positive number is treated as 1). The Dice similarity puts more weight on joint occurrences than on mismatches.

When comparing two rows, a match is counted for all columns with presences in both rows. Using M for the number of matches and N for the total number of columns with presence in just one row, we have

$$d_{jk} = 2M / (2M+N).$$

Jaccard

A similarity index for binary data. With the same notation as given for Dice similarity above, we have

$$d_{jk} = M / (M+N).$$

Kulczynski

A similarity index for binary data. With the same notation as given for Dice similarity above (with N_1 and N_2 referring to the two rows), we have

$$d_{jk} = \frac{\frac{M}{M+N_1} + \frac{M}{M+N_2}}{2}.$$

Ochiai

A similarity index for binary data, comparable to the cosine similarity for other data types:

$$d_{jk} = \sqrt{\frac{M}{M+N_1} \frac{M}{M+N_2}}.$$

Simpson

The Simpson index is defined simply as M / N_{\min} , where N_{\min} is the smaller of the numbers of presences in the two rows. This index treats two rows as identical if one is a subset of the other, making it useful for fragmentary data.

Bray-Curtis

Bray-Curtis is becoming the standard distance index for abundance data in e.g. marine biology.

$$d_{jk} = \frac{\sum_i |x_{ji} - x_{ki}|}{\sum_i (x_{ji} + x_{ki})}.$$

Cosine

The inner product of abundances each normalised to unit norm, i.e. the cosine of the angle between the vectors.

$$d_{jk} = \frac{\sum_i x_{ji} x_{ki}}{\sqrt{\sum_i x_{ji}^2} \sqrt{\sum_i x_{ki}^2}}.$$

Morisita

For abundance data.

$$\lambda_1 = \frac{\sum_i x_{ji} (x_{ji} - 1)}{\sum_i x_{ji} \left(\sum_i x_{ji} - 1 \right)}$$

$$\lambda_2 = \frac{\sum_i x_{ki} (x_{ki} - 1)}{\sum_i x_{ki} \left(\sum_i x_{ki} - 1 \right)}$$

$$d_{jk} = \frac{2 \sum_i x_{ji} x_{ki}}{(\lambda_1 + \lambda_2) \sum_i x_{ji} \sum_i x_{ki}}.$$

Raup-Crick

Raup-Crick index for absence-presence data. This index (Raup & Crick 1979) uses a randomization (Monte Carlo) procedure, comparing the observed number of species occurring in both associations with the distribution of co-occurrences from 200 random replicates.

Horn

Horn's overlap index for abundance data (Horn 1966).

$$N_j = \sum_i x_{ji}$$

$$N_k = \sum_i x_{ki}$$

$$d_{jk} = \frac{\sum_i [(x_{ji} + x_{ki}) \ln(x_{ji} + x_{ki})] - \sum_i x_{ji} \ln x_{ji} - \sum_i x_{ki} \ln x_{ki}}{(N_j + N_k) \ln(N_j + N_k) - N_j \ln N_j - N_k \ln N_k}.$$

Hamming

Hamming distance for categorical data as coded with integers (or sequence data coded as CAGT). The Hamming distance is the number of differences (mismatches), so that the distance between (3,5,1,2) and (3,7,0,2) equals 2. In PAST, this is normalised to the range [0,1], which is known to geneticists as "p-distance".

Chord

Euclidean distance between normalized vectors. Commonly used for abundance data. Can be written as

$$d_{jk} = \sqrt{2 - 2 \frac{\sum_i x_{ji} x_{ki}}{\sqrt{\sum_i x_{ji}^2} \sqrt{\sum_i x_{ki}^2}}}.$$

Manhattan

The sum of differences in each variable:

$$d_{jk} = \sum_i |x_{ji} - x_{ki}|.$$

Jukes-Cantor

Distance measure for genetic sequence data (CAGT). Similar to p (or Hamming) distance, but takes into account probability of reversals:

$$d = -\frac{3}{4} \ln \left(1 - \frac{4}{3} p \right)$$

Kimura

The Kimura 2-parameter distance measure for genetic sequence data (CAGT). Similar to Jukes-Cantor distance, but takes into account different probabilities of nucleotide transitions vs. transversions (Kimura 1980). With P the observed proportion of transitions and Q the observed number of transversions, we have

$$d = -\frac{1}{2} \ln(1 - 2P - Q) - \frac{1}{4} \ln(1 - 2Q).$$

Tajima-Nei

Distance measure for genetic sequence data (CAGT). Similar to Jukes-Cantor distance, but does not assume equal nucleotide frequencies.

User-defined similarity

Expects a symmetric similarity matrix rather than original data. No error checking!

User-defined distance

Expects a symmetric distance matrix rather than original data. No error checking!

Mixed

This option requires that data types have been assigned to columns (see *Entering and manipulating data*). A pop-up window will ask for the similarity/distance measure to use for each datatype. These will be combined using an average weighted by the number of variates of each type. The default choices correspond to those suggested by Gower, but other combinations may well work better. The "Gower" option is a range-normalised Manhattan distance.

All-zeros rows: Some similarity measures (Dice, Jaccard, Simpson etc.) are undefined when comparing two all-zero rows. To avoid errors, especially when bootstrapping sparse data sets, the similarity is set to zero in such cases.

Missing data: Most of these measures treat missing data (coded as '?') by pairwise deletion, meaning that if a value is missing in one of the variables in a pair of rows, that variable is omitted from the computation of the distance between those two rows. The exceptions are rho distance, using column average substitution, and Raup-Crick, which does not accept missing data. CHECK

References

Horn, H.S. 1966. Measurement of overlap in comparative ecological studies. *American Naturalist* 100:419-424.

Kimura, M. 1980. A simple model for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences. *Journal of Molecular Evolution* 16:111-120.

Raup, D. & R.E. Crick. 1979. Measurement of faunal similarity in paleontology. *Journal of Paleontology* 53:1213-1227.

Correlation

A matrix is presented with the correlations between all pairs of columns. Correlation values are given in the lower triangle of the matrix, and the probabilities that the columns are uncorrelated are given in the upper. Both parametric (Pearson) and non-parametric (Spearman and Kendall) coefficients and tests are available. Algorithms follow Press et al. (1992).

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

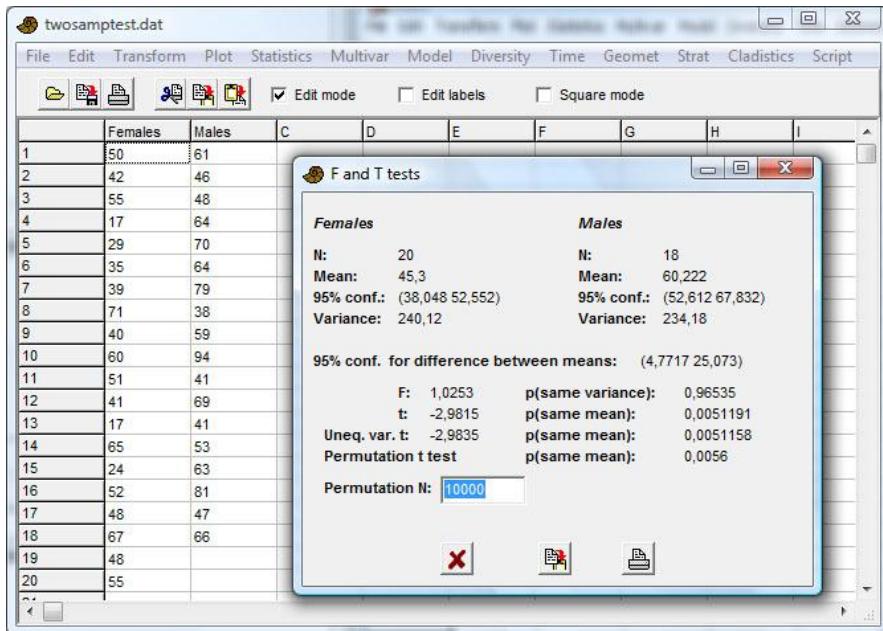
Var-covar

A symmetric matrix is presented with the variances and covariances between all pairs of columns.

Missing data: Supported by pairwise deletion.

F and t tests (two samples)

A number of classical, parametric statistics and tests for comparing the means and variances of two univariate samples (in two columns). Normal distribution is assumed.



Sample statistics

Means and variances are estimated as described above under Univariate statistics. The 95% confidence interval for the mean is based on the standard error for the estimate of the mean, and the t distribution. With s the estimate of the standard deviation, the confidence interval is

$$\left[\bar{x} - t_{(\alpha/2, n-1)} \frac{s}{\sqrt{n}}, \quad \bar{x} + t_{(\alpha/2, n-1)} \frac{s}{\sqrt{n}} \right].$$

Here, t has $n-1$ degrees of freedom, and $1-\alpha = 0.95$ for a 95% confidence interval.

The 95% confidence interval for the difference between the means accepts unequal sample sizes:

$$\left[|\bar{x} - \bar{y}| - t_{(\alpha/2, df)} s_D, \quad |\bar{x} - \bar{y}| + t_{(\alpha/2, df)} s_D \right]$$

where

$$SSE = \sum (x_i - \bar{x})^2 + \sum (y_i - \bar{y})^2$$

$$df = (n_1 - 1) + (n_2 - 1)$$

$$MSE = SSE / df$$

$$n_h = \frac{2}{1/n_1 + 1/n_2}$$

$$s_D = \sqrt{\frac{2MSE}{n_h}}.$$

F test

The *F* test has null hypothesis

H_0 : The two samples are taken from populations with equal variance.

The *F* statistic is the ratio of the larger variance to the smaller. The significance is two-tailed, with n_1 and n_2 degrees of freedom.

t test

The *t* test has null hypothesis

H_0 : The two samples are taken from populations with equal means.

From the standard error s_D of the difference of the means given above, the test statistic is

$$t = \frac{\bar{x} - \bar{y}}{s_D}.$$

Unequal variance *t* test

The unequal variance *t* test is also known as the Welch test. It can be used as an alternative to the basic *t* test when variances are very different, although it can be argued that testing for difference in the means in this case is questionable. The test statistic is

$$t = \frac{\bar{x} - \bar{y}}{\sqrt{\text{Var}(x)/n_1 + \text{Var}(y)/n_2}}.$$

The number of degrees of freedom is

$$df = \frac{\left[\frac{\text{Var}(x)}{n_1} + \frac{\text{Var}(y)}{n_2} \right]^2}{\frac{[\text{Var}(x)/n_1]^2}{n_1 - 1} + \frac{[\text{Var}(y)/n_2]^2}{n_2 - 1}}.$$

Permutation t test

The permutation test for equality of means uses the t statistic, but is non-parametric test with few assumptions. The number of permutations can be set by the user. The power of the test is limited by the sample size – significance at the $p<0.05$ level can only be achieved for $n>3$ in each sample.

T test (one sample)

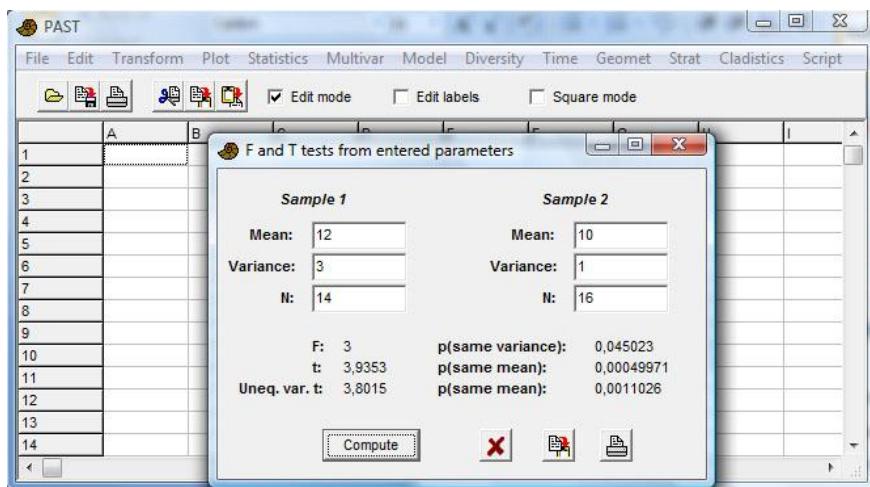
The one-sample t test is used to investigate whether the sample is likely to have been taken from a population with a given (theoretical) mean.

The 95% confidence interval for the mean is calculated using the t distribution.

Missing data: Supported by deletion.

F and t tests from parameters

Sometimes publications give not the data, but values for sample size, mean and variance for two samples. These can be entered manually using the 'F and t from parameters' option in the menu. This module does not use any data from the spreadsheet.

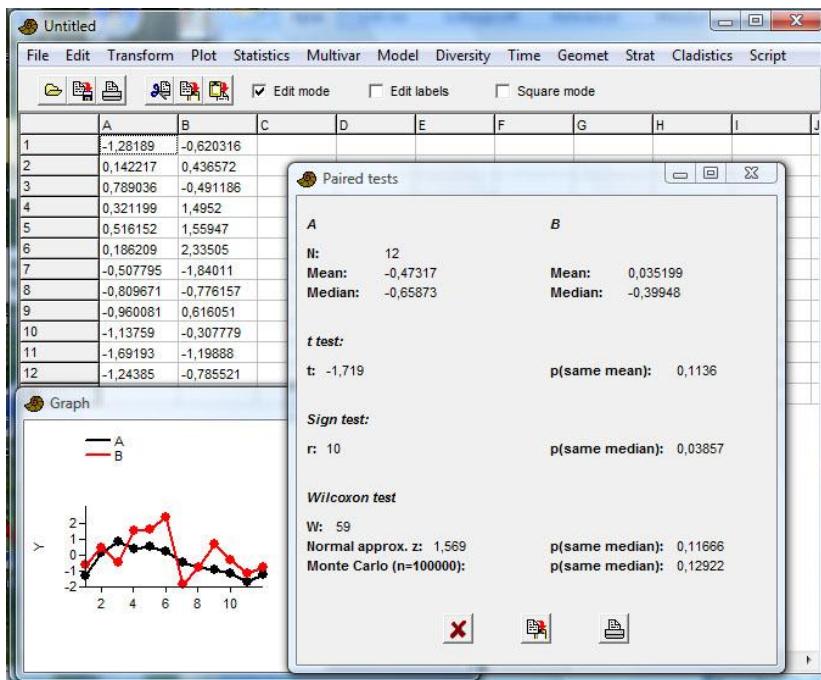


Paired tests (*t*, sign, Wilcoxon)

Three statistical tests (one parametric, two non-parametric) for two samples (columns) of univariate data. The data points are paired, meaning that the two values in each row are associated. For example, the test could be the for length of the left vs. the right arm in a number of people, or the diversity in summer vs. winter at a number of sites. Controlling for a “nuisance factor” (person, site) in this way increases the power of the test. The null hypothesis is:

H_0 : The mean (*t* test) or median (sign test, Wilcoxon test) of the difference is zero.

All reported *p* values are two-tailed.



t test

Testing for mean difference equal to zero using the normal one-sample *t* test. With $d_i = x_i - y_i$, we have

$$s = \sqrt{\frac{1}{n-1} \sum (d_i - \bar{d})^2},$$

$$t = \frac{\bar{d}}{s/\sqrt{n}}.$$

There are $n-1$ degrees of freedom. This test assumes normal distribution of the differences.

Sign test

The sign (binomial) test simply counts the number of cases n_1 where $x_i > y_i$ and n_2 where $y_i > x_i$. The number $\max(n_1, n_2)$ is reported. The p value is exact, computed from the binomial distribution. The sign test will typically have lower power than the other paired tests, but make few assumptions.

Wilcoxon signed rank test

A non-parametric rank test that does not assume normal distribution. The null hypothesis is no median shift (no difference).

All rows with zero difference are first removed by the program. Then the absolute values of the differences $|d_i|$ are ranked (R_i), with mean ranks assigned for ties. The sum of ranks for pairs where d_i is positive is W^+ . The sum of ranks for pairs where d_i is negative is W . The reported test statistic is

$$W = \max(W^+, W^-)$$

(note that there are several other, equivalent versions of this test, reporting other statistics).

For large n (say $n > 10$), the large-sample approximation to p can be used. This depends on the normal distribution of the test statistic W :

$$E(W) = \frac{n(n+1)}{4}$$
$$Var(W) = \frac{n(n+1)(2n+1)}{24} - \frac{\sum_g f_g^3 - f_g}{48}.$$

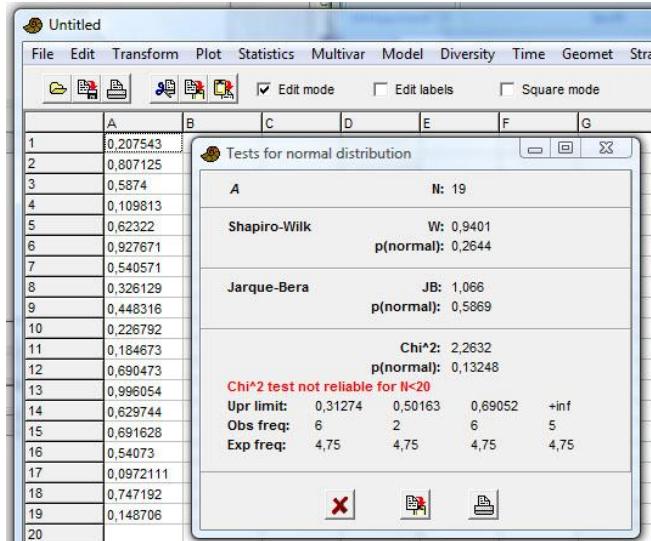
The last term is a correction for ties, where f_g is the number of elements in tie g . The resulting z is reported, together with the p value.

For smaller sample sizes (or in general), the Monte Carlo significance value should be used instead. It is based on 100,000 random reassessments of values to columns, within each pair. This value will be practically identical to the exact p value (which may be included in a future version).

Missing data: Supported by deletion of the row.

Normality (one sample)

Three statistical tests for normal distribution of one sample of univariate data, given in one column. The data below were generated by a random number generator with uniform distribution.



For all the three tests, the null hypothesis is

H_0 : The sample was taken from a population with normal distribution.

If the given $p(\text{normal})$ is less than 0.05, normal distribution can be rejected. Of the three given tests, the Shapiro-Wilk is considered to be the more exact, and the two other tests (Jarque-Bera and a chi-square test) are given for reference. There is a maximum sample size of $n=5000$, while the minimum sample size is 3 (the tests will of course have extremely small power for such small n).

Shapiro-Wilk test

The Shapiro-Wilk test (Shapiro & Wilk 1965) returns a test statistic W , which is small for non-normal samples, and a p value. The implementation is based on the standard code "AS R94" (Royston 1995), correcting an inaccuracy in the previous algorithm "AS 181" for large sample sizes.

Jarque-Bera test

The Jarque-Bera test (Jarque & Bera 1987) is based on skewness S and kurtosis K . The test statistic is

$$JB = \frac{n}{6} \left(S^2 + \frac{(K - 3)^2}{4} \right)$$

In this context, the skewness and kurtosis used are

$$S = \frac{1}{n} \frac{\sum (x_i - \bar{x})^3}{\left(\sqrt{\frac{1}{n} \sum (x_i - \bar{x})^2} \right)^3},$$

$$K = \frac{1}{n} \frac{\sum (x_i - \bar{x})^4}{\left(\sqrt{\frac{1}{n} \sum (x_i - \bar{x})^2} \right)^4}.$$

Note that these equations contain simpler estimators than the G_1 and G_2 given above, and that the kurtosis here will be 3, not zero, for a normal distribution.

Asymptotically (for large sample sizes), the test statistic has a chi-square distribution with two degrees of freedom, and this forms the basis for the p value given by Past. It is known that this approach works well only for large sample sizes, and a future version might include a method based on Monte Carlo simulation.

Chi-square test

The chi-square test uses an expected normal distribution in four bins, based on the mean and standard deviation estimated from the sample, and constructed to have equal expected frequencies in all bins. The upper limits of all bins, and the observed and expected frequencies, are displayed. A warning message is given if $n < 20$, i.e. expected frequency less than 5 in each bin. There is 1 degree of freedom. This test is both theoretically questionable and has low power, and is not recommended. It is included for reference.

Missing data: Supported by deletion.

References

Jarque, C. M. & Bera, A. K. 1987. A test for normality of observations and regression residuals. *International Statistical Review* 55:163–172.

Royston, P. 1995. A remark on AS 181: The W -test for normality. *Applied Statistics* 44:547-551.

Shapiro, S. S. & Wilk, M. B. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52:591–611.

Chi²

The Chi-square test is the one to use if your data consist of the numbers of elements in different bins (compartments). For example, this test can be used to compare two associations (columns) with the number of individuals in each taxon organized in the rows. You should be a little cautious about such comparisons if any of the bins contain less than five individuals.

There are two options that you should select or not for correct results. "Sample vs. expected" should be ticked if your second column consists of values from a theoretical distribution (expected values) with zero error bars. If your data are from two counted samples each with error bars, leave this box open. This is *not* a small-sample correction.

"One constraint" should be ticked if your expected values have been normalized in order to fit the total observed number of events, or if two counted samples necessarily have the same totals (for example because they are percentages). This will reduce the number of degrees of freedom by one.

When "one constraint" is selected, a permutation test is available, with 10000 random replicates. For "Sample vs. expected" these replicates are generated by keeping the expected values fixed, while the values in the first column are random with relative probabilities as specified by the expected values, and with constant sum. For two samples, all cells are random but with constant row and column sums.

With one constraint and a 2x2 table, the Fisher's exact test is also given (two-tailed). When available, the Fisher's exact test is far superior to the chi-square.

See e.g. Brown & Rothery (1993) or Davis (1986) for details.

Missing data: Supported by deletion.

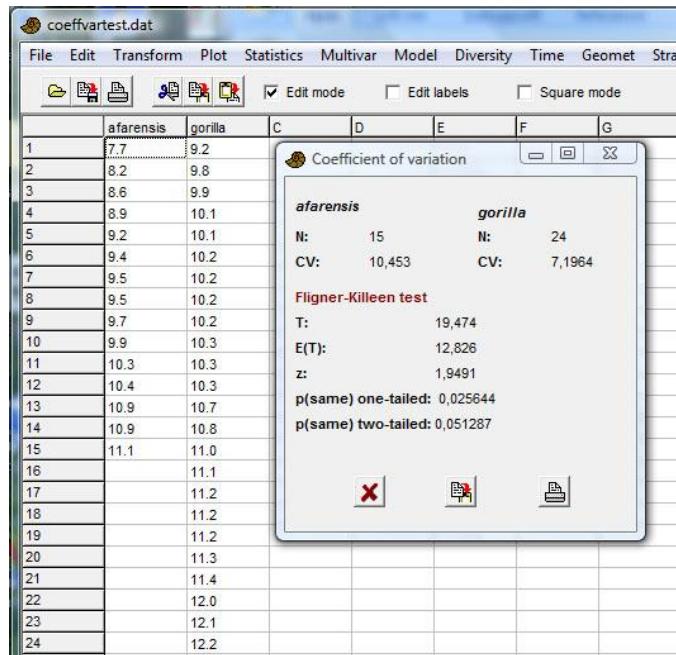
References

Brown, D. & P. Rothery. 1993. Models in biology: mathematics, statistics and computing. John Wiley & Sons.

J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Coefficient of variation

This module tests for equal coefficient of variation in two samples, given in two columns.



The coefficient of variation (or relative variation) is defined as the ratio of standard deviation to the mean in percent, and is computed as:

$$CV = \frac{s}{\bar{x}} \cdot 100 = \frac{\sqrt{\frac{1}{n-1} \sum (x_i - \bar{x})^2}}{\bar{x}} \cdot 100$$

The null hypothesis if the statistical test is:

H_0 : The samples were taken from populations with the same coefficient of variation.

If the given $p(\text{normal})$ is less than 0.05, normal distribution can be rejected. Donnelly & Kramer (1999) describe the coefficient of variation and review a number of statistical tests for the comparison of two samples. They recommend the Fligner-Killeen test (Fligner & Killeen 1976), as implemented in Past. This test is both powerful and is relatively insensitive to distribution. The following statistics are reported:

T : The Fligner-Killeen test statistic, which is a sum of transformed ranked positions of the smaller sample within the pooled sample (see Donnelly & Kramer 1999 for details).

$E(T)$: The expected value for T .

z : The z statistic, based on T , $\text{Var}(T)$ and $E(T)$. Note this is a large-sample approximation.

p: The $p(H_0)$ value. Both the one-tailed and two-tailed values are given. For the alternative hypothesis of difference in either direction, the two-tailed value should be used. However, the Fligner-Killeen test has been used to compare variation within a sample of fossils with variation within a closely related modern species, to test for multiple fossil species (Donnelly & Kramer 1999). In this case the alternative hypothesis might be that CV is larger in the fossil population, if so then a one-tailed test can be used for increased power.

The screenshot above reproduces the example of Donnelly & Kramer (1999), showing that the relative variation within *Australopithecus afarensis* is significantly larger than in *Gorilla gorilla*. This could indicate that *A. afarensis* represents several species.

References

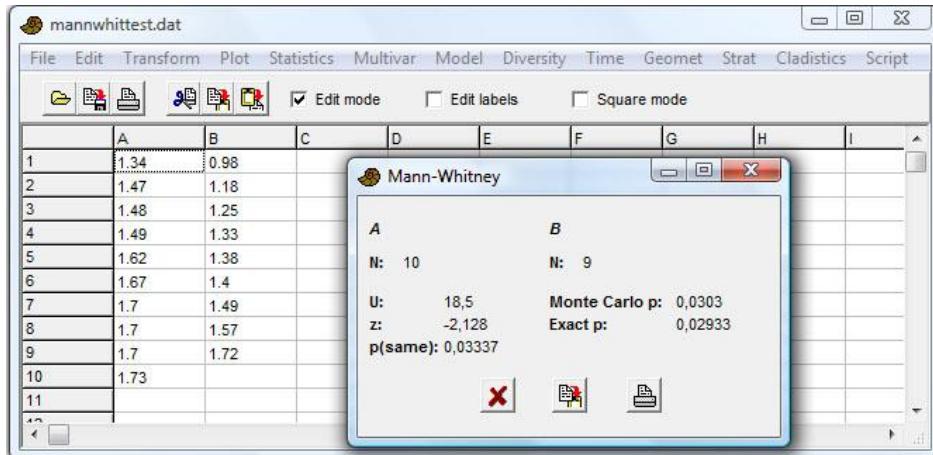
Donnelly, S.M. & Kramer, A. 1999. Testing for multiple species in fossil samples: An evaluation and comparison of tests for equal relative variation. *American Journal of Physical Anthropology* 108:507-529.

Fligner, M.A. & Killeen, T.J. 1976. Distribution-free two sample tests for scale. *Journal of the American Statistical Association* 71:210-213.

Mann-Whitney test

The two-tailed (Wilcoxon) Mann-Whitney U test can be used to test whether the medians of two independent samples are different. It does not assume normal distribution, but does assume equal-shaped distribution in both groups. The null hypothesis is

H_0 : The two samples are taken from populations with equal medians.



This test is non-parametric, which means that the distributions can be of any shape.

For each value in sample 1, count number of values in sample 2 that are smaller than it (ties count 0.5). The total of these counts is the test statistic U (sometimes called T). If the value of U is smaller when reversing the order of samples, this value is chosen instead (it can be shown that $U_1+U_2=n_1n_2$).

In the left column is given an asymptotic approximation to p based on the normal distribution (two-tailed), which is only valid for large n . It includes a continuity correction and a correction for ties:

$$z = \frac{U - n_1 n_2 / 2 + 0.5}{\sqrt{\frac{n_1 n_2 \left(n^3 - n - \sum_g f_g^3 - f_g \right)}{12n(n-1)}}$$

where $n=n_1+n_2$ and f_g is the number of elements in tie g .

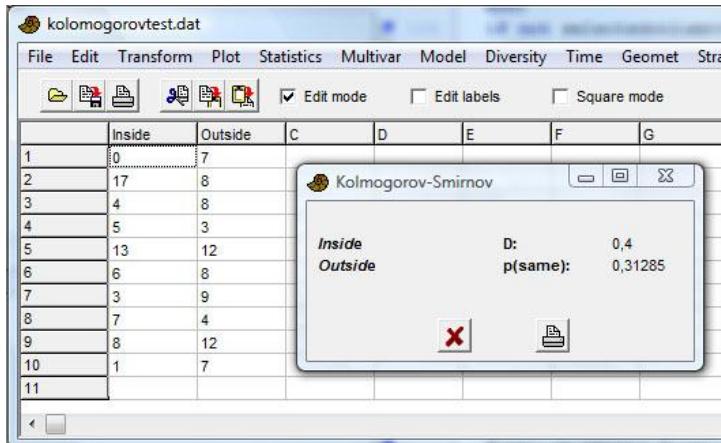
For $n_1+n_2 \leq 30$ (e.g. 15 values in each group), an exact p value is given, based on all possible combinations of group assignment. If available, always use this exact value. For larger samples, the asymptotic approximation is quite accurate. A Monte Carlo value based on 10 000 random assignments is also given – the purpose of this is mainly as a control on the asymptotic value.

Missing data: Supported by deletion.

Kolmogorov-Smirnov

The Kolmogorov-Smirnov test is a nonparametric test for overall equal distribution of two univariate samples. In other words, it does not test specifically for equality of mean, variance or any other parameter. The null hypothesis is

H_0 : The two samples are taken from populations with equal distribution.



In the version of the test provided by Past, both columns must represent samples. You can not test a sample against a theoretical distribution (one-sample test).

The test statistic is the maximum absolute difference between the two empirical cumulative distribution functions:

$$D = \max_x |S_{N_1}(x) - S_{N_2}(x)|$$

The algorithm is based on Press et al. (1992), with significance estimated after Stephens (1970).

Define the function

$$Q_{KS}(\lambda) = 2 \sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2\lambda^2}$$

With $N_e = N_1 N_2 / (N_1 + N_2)$, the significance is computed as

$$p = Q_{KS} \left(\left[\sqrt{N_e} + 0.12 + 0.11 / \sqrt{N_e} \right] D \right)$$

Missing data: Supported by deletion.

References

Press, W.H., Teukolsky, S.A., Vetterling, W.T. & Flannery, B.P. 1992. Numerical Recipes in C. 2nd edition. Cambridge University Press.

Stephens, M.A. 1970. Use of the Kolmogorov-Smirnov, Cramer-von Mises and related statistics without extensive tables. *Journal of the Royal Statistical Society, Series B* 32:115-122.

Spearmann/Kendall

These non-parametric rank-order tests are used to test for correlation between two variables. The permutation test is based on 1000 random replicates. Algorithms after Press et al. (1992),

These statistics are also available through the “Correlation” module, but then without the permutation option.

Missing data: Supported by deletion.

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

Contingency table

A contingency table is input to this routine. Rows represent the different states of one nominal variable, columns represent the states of another nominal variable, and cells contain the counts of occurrences of that specific state (row, column) of the two variables. A measure and probability of association of the two variables (based on Chi-square) is then given.

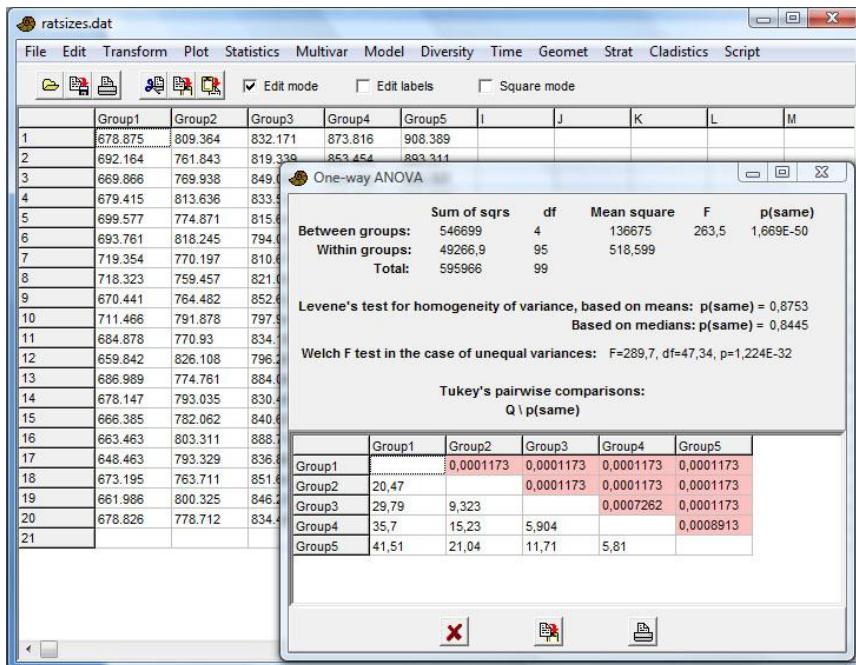
For example, rows may represent taxa and columns samples as usual (with specimen counts in the cells). The contingency table analysis then gives information on whether the two variables of taxon and locality are associated. If not, the data matrix is not very informative. For details, see Press et al. (1992).

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

One-way ANOVA

One-way ANOVA (analysis of variance) is a statistical procedure for testing the null hypothesis that several univariate samples (in columns) are taken from populations with the same mean. The samples are assumed to be close to normally distributed and have similar variances. If the sample sizes are equal, these two assumptions are not critical. If the assumptions are strongly violated, the nonparametric Kruskal-Wallis test should be used instead.



ANOVA table

The between-groups sum of squares is given by:

$$SS_{bg} = \sum_g n_g (\bar{x}_g - \bar{x}_T)^2$$

where n_g is the size of group g , and the means are group and total means. The between-groups sum of squares has an associated df_{bg} , the number of groups minus one.

The within-groups sum of squares is

$$SS_{wg} = \sum_g \sum_i (x_i - \bar{x}_g)^2$$

where the x_i are those in group g . The within-groups sum of square has an associated df_{wg} , the total number of values minus the number of groups.

The mean squares between and within groups are given by

$$MS_{bg} = \frac{SS_{bg}}{df_{bg}}$$

$$MS_{wg} = \frac{SS_{wg}}{df_{wg}}$$

Finally, the test statistic F is computed as

$$F = \frac{MS_{bg}}{MS_{wg}}$$

The p value is based on F with df_{bg} and df_{wg} degrees of freedom.

Levene's test

Levene's test for homogeneity of variance (homoskedasticity), that is, whether variances are equal as assumed by ANOVA, is also given. Two versions of the test are included. The original Levene's test is based on means. This version has more power if the distributions are normal or at least symmetric. The version based on medians has less power, but is more robust to non-normal distributions. Note that this test can be used also for only two samples, giving an alternative to the F test for two samples described above.

Unequal-variance (Welch) ANOVA

If Levene's test is significant, meaning that you have unequal variances, you can use the unequal-variance (Welch) version of ANOVA, with the F , df and p values given.

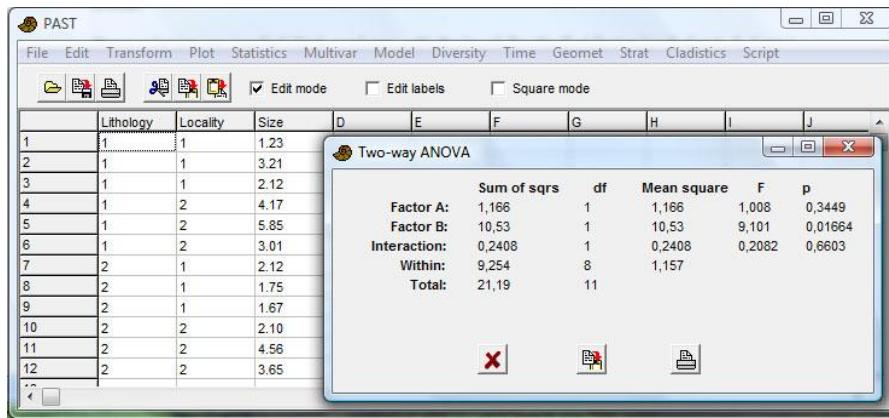
Post-hoc pairwise tests

If the ANOVA shows significant inequality of the means (small p), you can go on to study the given table of "post-hoc" pairwise comparisons, based on Tukey's HSD (Honestly Significant Difference) test. The Studentized Range Statistic Q is given in the lower left triangle of the array, and the probabilities $p(equal)$ in the upper right. Sample sizes do not have to be equal for the version of Tukey's test used.

Two-way ANOVA

Two-way ANOVA (analysis of variance) is a statistical procedure for testing the null hypotheses that several univariate samples have the same mean across each of the two factors, and that there are no dependencies (interactions) between factors. The samples are assumed to be close to normally distributed and have similar variances. If the sample sizes are equal, these two assumptions are not critical.

Three columns are needed. First, a column with the levels for the first factor (coded as 1, 2, 3 etc.), then a column with the levels for the second factor, and finally the column of the corresponding measured values.

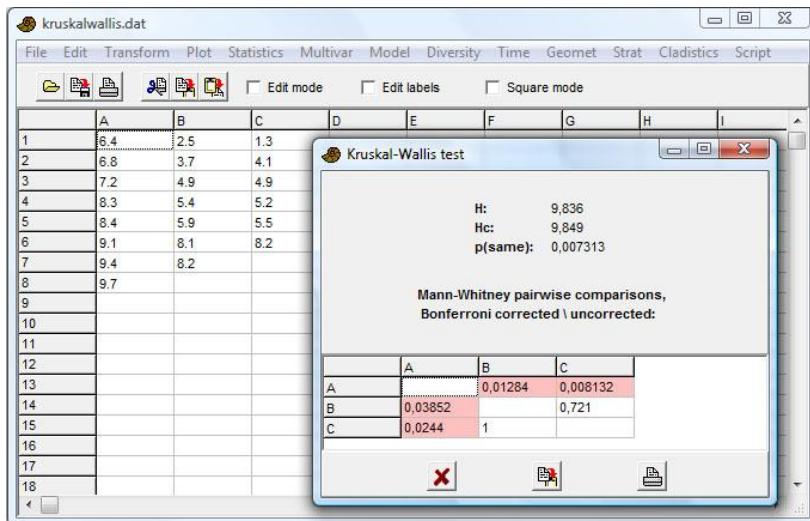


The algorithm uses weighted means for unbalanced designs.

Kruskal-Wallis

The Kruskal-Wallis test is a non-parametric ANOVA, comparing the medians of several univariate groups (given in columns). It can also be regarded as a multiple-group extension of the Mann-Whitney test (Zar 1996). It does not assume normal distribution, but does assume equal-shaped distribution for all groups. The null hypothesis is

H_0 : The samples are taken from populations with equal medians.



The test statistic H is computed as follows:

$$H = \frac{12}{n(n+1)} \left(\sum_g \frac{T_g^2}{n_g} \right) - 3(n+1),$$

where n_g is the number of elements in group g , n is the total number of elements, and T_g is the sum of ranks in group g .

The test statistic H_c is adjusted for ties:

$$H_c = \frac{H}{\sum_i f_i^3 - f_i} \cdot \frac{1 - \frac{i}{n^3 - n}}{1 - \frac{i}{n^3 - n}},$$

where f_i is the number of elements in tie i .

With G the number of groups, the p value is approximated from H_c using the chi-square distribution with $G-1$ degrees of freedom. This is less accurate if any $n_g < 5$.

Post-hoc pairwise tests

Mann-Whitney pairwise test p values are given for all $N_p=G(G-1)/2$ pairs of groups, in the upper right triangle of the matrix. The lower right triangle gives the corresponding p values, but multiplied with N_p as a conservative correction for multiple testing (Bonferroni correction). The values use the asymptotic approximation described under the Mann-Whitney module. If samples are very small, it may be useful to run the exact test available in that module instead.

Missing data: Supported by deletion.

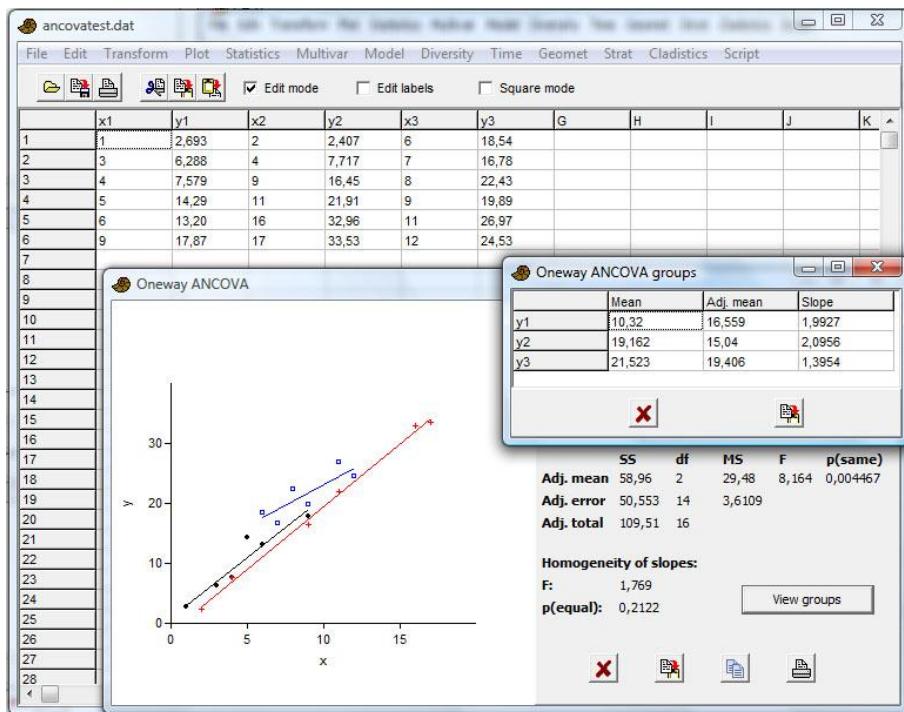
Reference

Zar, J.H. 1996. Biostatistical analysis. 3rd ed. Prentice Hall.

One-way ANCOVA

ANCOVA (Analysis of covariance) tests for equality of means for several univariate groups, adjusted for covariance with another variate. ANCOVA can be compared with ANOVA, but has the added feature that for each group, variance that can be explained by a specified "nuisance" covariate (x) is removed. This adjustment can increase the power of the test substantially.

The program expects two or more pairs of columns, where each pair (group) is a set of correlated x-y data (means are compared for y, while x is the covariate). The example below uses three pairs (groups).



The program presents a scatter plot and linear regression lines for all the groups. The ANOVA-like summary table contains sum-of-squares etc. for the adjusted means (between-groups effect) and adjusted error (within-groups), together with an *F* test for the adjusted means. An *F* test for the equality of regression slopes (as assumed by the ANCOVA) is also given. In the example, equal adjusted means in the three groups can be rejected at $p<0.005$. Equality of slopes can not be rejected ($p=0.21$).

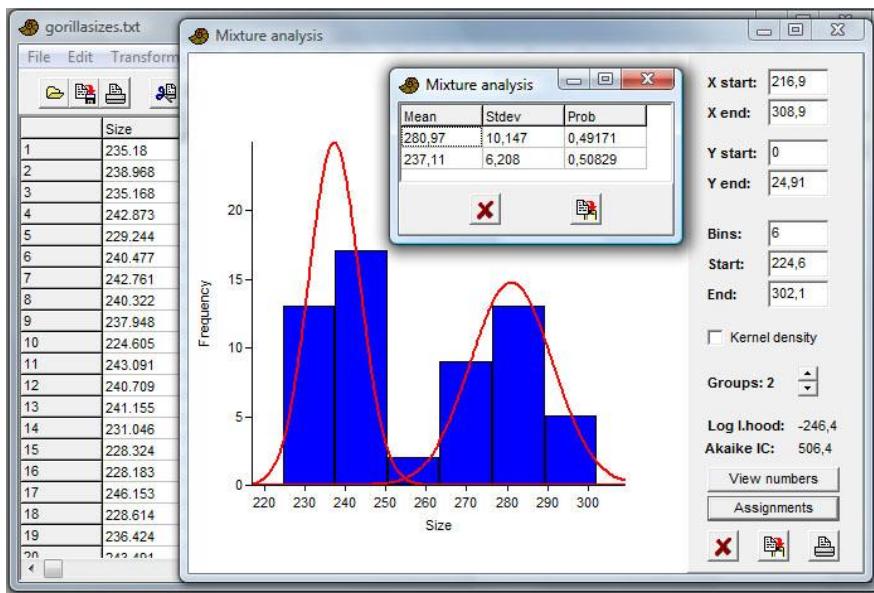
"View groups" gives the summary statistics for each group (mean, adjusted mean and regression slope).

Assumptions include similar linear regression slopes for all groups, normal distributions, similar variance and sample sizes.

Mixture analysis

Mixture analysis is a maximum-likelihood method for estimating the parameters (mean, standard deviation and proportion) of two or more univariate normal distributions, based on a pooled univariate sample. For example, the method can be used to study differences between sexes (two groups), or several species, or size classes, when no independent information about group membership is available.

The program expects one column of univariate data, assumed to be taken from a mixture of normally distributed populations. In the example below, sizes of male and female gorillas have been pooled in one sample. The means, standard deviations and proportions of the two original samples have been almost perfectly recovered (see "Univariate" above).



PAST uses the EM algorithm (Dempster et al. 1977), which can get stuck on a local optimum. The procedure is therefore automatically run 20 times, each time with new, random starting positions for the means. The starting values for standard deviation are set to s/G , where s is the pooled standard deviation and G is the number of groups. The starting values for proportions are set to $1/G$. The user is still recommended to run the program a few times to check for stability of the solution ("better" solutions have less negative log likelihood values).

The Akaike Information Criterion (AIC; Akaike 1974) is calculated with a small-sample correction:

$$AICc = 2k - 2 \ln L + \frac{2k(k+1)}{n-k-1}$$

where k is the number of parameters, n the number of data points and L the likelihood of the model given the data. A minimal value for AIC indicates that you have chosen the number of groups that produces the best fit without overfitting.

It is possible to assign each of the data points to one of the groups with a maximum likelihood approach. This can be used as a non-hierarchical clustering method for univariate data. The

“Assignments” button will open a window where the value of each Gaussian probability density function is given for each data point. The data point can be assigned to the group that shows the largest value.

Missing data: Supported by deletion.

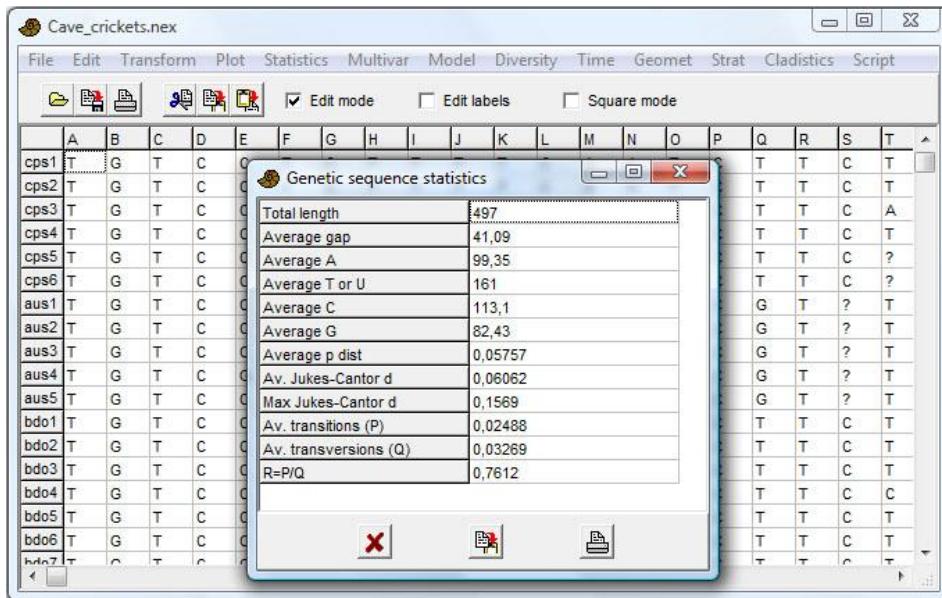
References

Akaike, H. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19: 716-723.

Dempster, A.P., Laird, N.M. & Rubin, D.B. 1977. Maximum likelihood from incomplete data via the EM algorithm". *Journal of the Royal Statistical Society, Series B* 39:1-38.

Genetic sequence stats

A number of simple statistics on genetic sequence (DNA or RNA) data. The module expects a number of rows, each with a sequence. The sequences are expected to be aligned and of equal length including gaps (coded as '?'). Some of these statistics are useful for selecting appropriate distance measures elsewhere in Past.



Total length: The total sequence length, including gaps, of one sequence

Average gap: The number of gap positions, averaged over all sequences

Average A, T/U, C, G: The average number of positions containing each nucleotide

Average *p* distance: The *p* distance between two sequences, averaged over all pairs of sequences. The *p* (or Hamming) distance is defined as the proportion of unequal positions

Average Jukes-Cantor *d*: The Jukes-Cantor *d* distance between two sequences, averaged over all pairs of sequences. $d = -3\ln(1 - 4p/3)/4$, where *p* is the *p* distance

Maximal Jukes-Cantor *d*: Maximal Jukes-Cantor distance between any two sequences

Average transitions (P): Average number of transitions (a↔g, c↔t, i.e. within purines, pyrimidines)

Average transversions (Q): Average number of transversions (a↔t, a↔c, c↔g, t↔g, i.e. across purines, pyrimidines)

R=P/Q: The transition/transversion ratio

Missing data: Treated as gaps.

Multivar menu

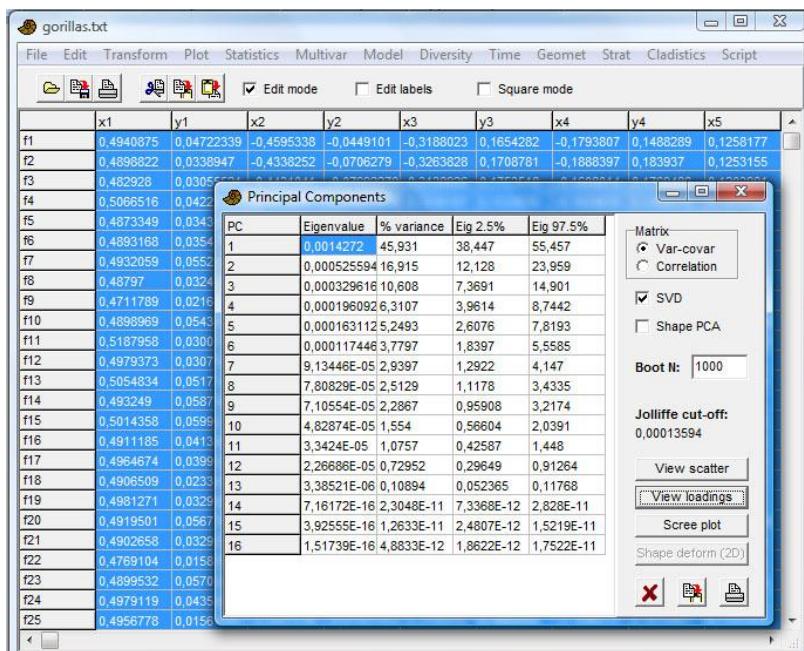
Principal components

Principal components analysis (PCA) is a procedure for finding hypothetical variables (components) which account for as much of the variance in your multidimensional data as possible (Davis 1986, Harper 1999). These new variables are linear combinations of the original variables. PCA may be used for simple reduction of the data set to only two variables (the two most important components), for plotting purposes. One might also try to hypothesize that the most important components are correlated with some other underlying variables. For morphometric data, this might be size, while for ecological data it might be a physical gradient (e.g. temperature or depth).

The input data consists of a matrix of multivariate data, with items in rows and variates in columns. There is no separate centering of groups before eigenanalysis – the analysis does not take groups into account.

The PCA routine finds the eigenvalues and eigenvectors of the variance-covariance matrix or the correlation matrix. Choose var-covar if all your variables are measured in the same units (e.g. centimetres). Choose correlation (normalized var-covar) if your variables are measured in different units; this implies normalizing all variables using division by their standard deviations. The eigenvalues, giving a measure of the variance accounted for by the corresponding eigenvectors (components) are given for all components. The percentages of variance accounted for by these components are also given. If most of the variance is accounted for by the first one or two components, you have scored a success, but if the variance is spread more or less evenly among the components, the PCA has in a sense not been very successful.

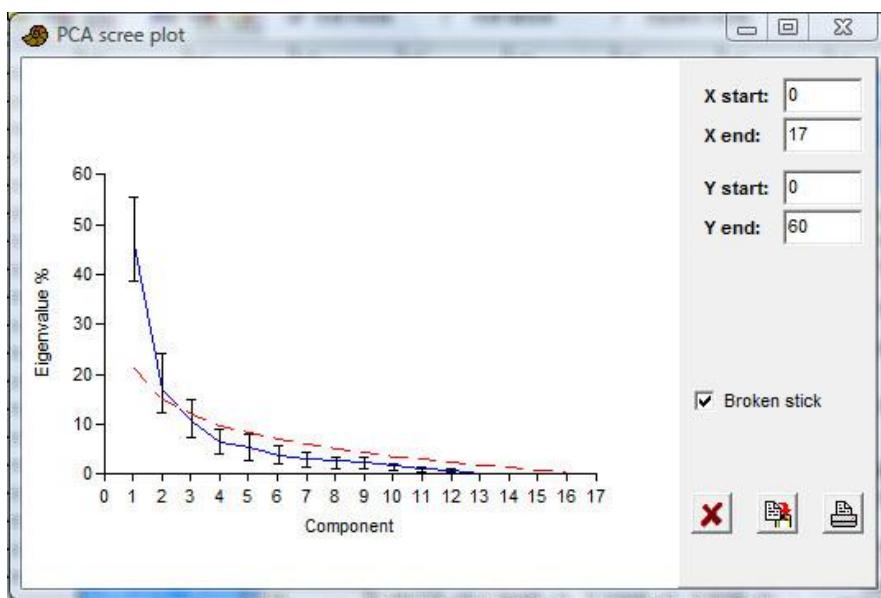
In the example below (landmarks from gorilla skulls), component 1 is strong, explaining 45.9% of variance. The bootstrapped confidence intervals (see below) are not shown unless the 'Boot N' value is non-zero.



The Jolliffe cut-off value gives an informal indication of how many principal components should be considered significant (Jolliffe, 1986). Components with eigenvalues smaller than the Jolliffe cut-off may be considered insignificant, but too much weight should not be put on this criterion.

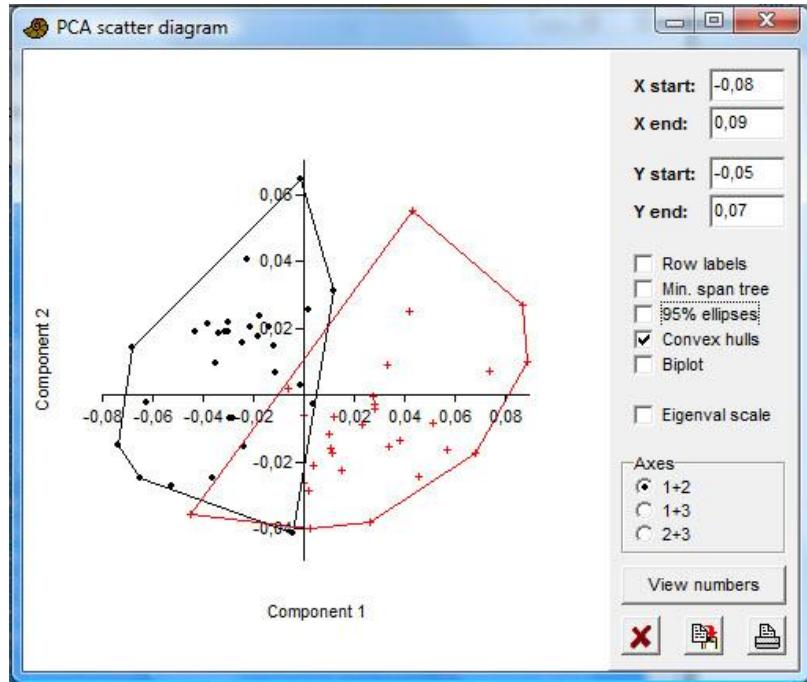
Row-wise bootstrapping is carried out if a non-zero number of bootstrap replicates (e.g. 1000) is given in the 'Boot N' box. The bootstrapped components are re-ordered and reversed according to Peres-Neto et al. (2003) to increase correspondence with the original axes. 95% bootstrapped confidence intervals are given for the eigenvalues.

The 'Scree plot' (simple plot of eigenvalues) can also be used to indicate the number of significant components. After this curve starts to flatten out, the corresponding components may be regarded as insignificant. 95% confidence intervals are shown if bootstrapping has been carried out. The eigenvalues expected under a random model (Broken Stick) are optionally plotted - eigenvalues under this curve may represent non-significant components (Jackson 1993).



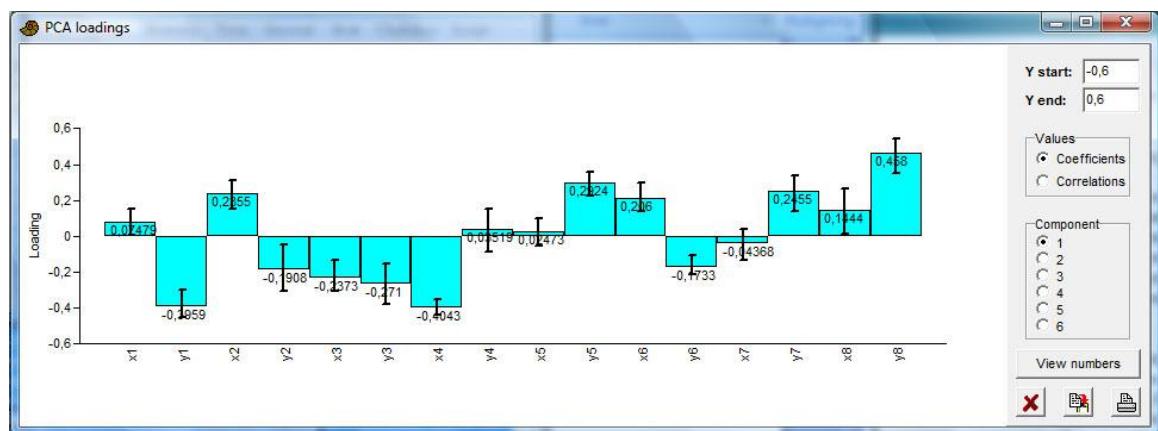
In the gorilla example above, the eigenvalues for the 16 components (blue line) lie above the broken stick values (red dashed line) for the first two components, although the broken stick is inside the 95% confidence interval for the second component.

The 'View scatter' option allows you to see all your data points (rows) plotted in the coordinate system given by the two most important components. If you have colored (grouped) rows, the different groups will be shown using different symbols and colours. You can also plot the Minimal Spanning Tree, which is the shortest possible set of connected lines connecting all points. This may be used as a visual aid in grouping close points. The MST is based on an Euclidean distance measure of the original data points, so it is most meaningful when all your variables use the same unit. The 'Biplot' option will show a projection of the original axes (variables) onto the scattergram. This is another visualisation of the PCA loadings (coefficients) - see below.



If the "Eigenval scale" is ticked, the data points will be scaled by $1/\sqrt{d_k}$, and the biplot eigenvectors by $\sqrt{d_k}$ - this is the correlation biplot of Legendre & Legendre (1998). If not ticked, the data points are not scaled, while the biplot eigenvectors are normalized to equal length (but not to unity, for graphical reasons) - this is the distance biplot.

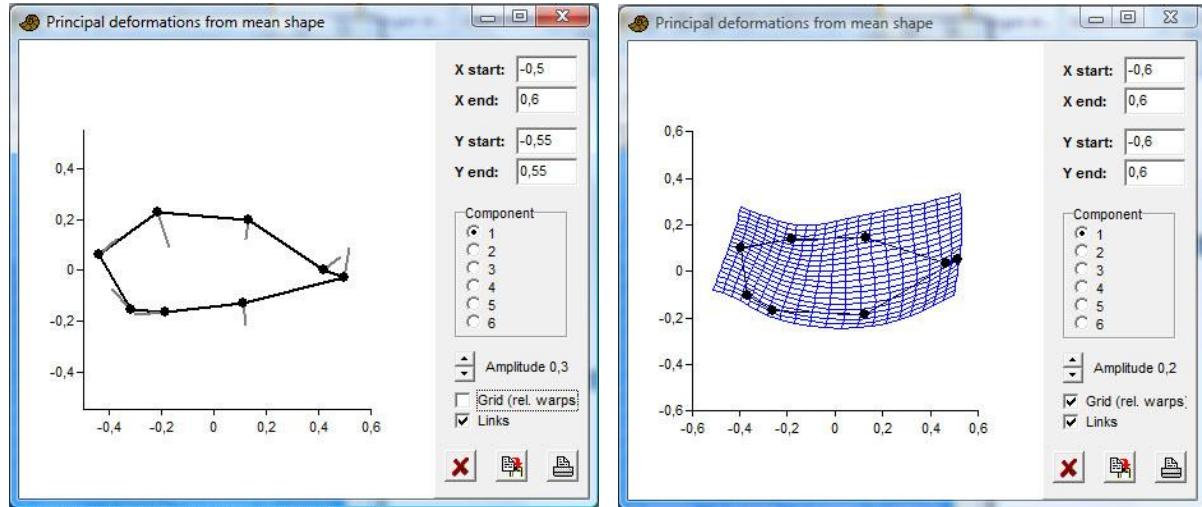
The 'View loadings' option shows to what degree your different original variables (given in the original order along the x axis) enter into the different components (as chosen in the radio button panel). These component loadings are important when you try to interpret the 'meaning' of the components. The 'Coefficients' option gives the PC coefficients, while 'Correlation' gives the correlation between a variable and the PC scores. If bootstrapping has been carried out, 95% confidence intervals are shown (only for the Coefficients option).



The 'SVD' option will enforce use of the superior Singular Value Decomposition algorithm instead of "classical" eigenanalysis. The two algorithms will normally give almost identical results, but axes may be flipped.

The 'Shape PCA' option enforces PCA on the var-covar matrix, and enables the 'Shape deform' button. This option is designed for 2D landmark position data. The default Shape Deformation plot is

a ‘lollipop plot’ with the mean shape shown as dots and vectors (lines) pointing in the directions of the axis loadings. The “Grid” option shows the thin-plate spline deformation grids corresponding to the different components. This is in effect a “relative warps” analysis, including the uniform component. For relative warps without the uniform component, see “Relative warps” in the Geometry menu.



Bruton & Owen (1988) describe a typical morphometrical application of PCA.

Missing data is supported by column average substitution.

References

Bruton, D.L. & A.W. Owen. 1988. The Norwegian Upper Ordovician illaenid trilobites. *Norsk Geologisk Tidsskrift* 68:241-258.

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Harper, D.A.T. (ed.). 1999. Numerical Palaeobiology. John Wiley & Sons.

Jackson, D.A. 1993. Stopping rules in principal components analysis: a comparison of heuristical and statistical approaches. *Ecology* 74:2204-2214.

Jolliffe, I.T. 1986. Principal Component Analysis. Springer-Verlag.

Peres-Neto, P.R., D.A. Jackson & K.M. Somers. 2003. Giving meaningful interpretation to ordination axes: assessing loading significance in principal component analysis. *Ecology* 84:2347-2363.

Principal coordinates

Principal coordinates analysis (PCO) is another ordination method, somewhat similar to PCA. It is also known as Metric Multidimensional Scaling (different from Non-metric Multidimensional Scaling!). The algorithm is taken from Davis (1986).

The PCO routine finds the eigenvalues and eigenvectors of a matrix containing the distances or similarities between all data points. The Gower measure will normally be used instead of Euclidean distance, which gives results similar to PCA. An additional eleven distance measures are available - these are explained under Cluster Analysis. The eigenvalues, giving a measure of the variance accounted for by the corresponding eigenvectors (coordinates) are given for the first four most important coordinates (or fewer if there are fewer than four data points). The percentages of variance accounted for by these components are also given.

The similarity/distance values are raised to the power of c (the "Transformation exponent") before eigenanalysis. The standard value is $c=2$. Higher values (4 or 6) may decrease the "horseshoe" effect (Podani & Miklos 2002).

The 'View scatter' option allows you to see all your data points (rows) plotted in the coordinate system given by the PCO. If you have colored (grouped) rows, the different groups will be shown using different symbols and colours. The "Eigenvalue scaling" option scales each axis using the square root of the eigenvalue (recommended). The minimal spanning tree option is based on the selected similarity or distance index in the original space.

Missing data is supported by pairwise deletion (not for the Raup-Crick, Rho or user-defined indices).

References

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Podani, J. & I. Miklos. 2002. Resemblance coefficients and the horseshoe effect in principal coordinates analysis. *Ecology* 83:3331-3343.

Non-metric MDS

Non-metric multidimensional scaling is based on a distance matrix computed with any of 13 supported distance measures, as explained under Cluster Analysis below. The algorithm then attempts to place the data points in a two- or three-dimensional coordinate system such that the *ranked differences* are preserved. For example, if the original distance between points 4 and 7 is the ninth largest of all distances between any two points, points 4 and 7 will ideally be placed such that their euclidean distance in the 2D plane or 3D space is still the ninth largest. Non-metric multidimensional scaling intentionally does not take absolute distances into account.

The program may converge on a different solution in each run, depending upon the random initial conditions. Each run is actually a sequence of 11 trials, from which the one with smallest stress is chosen. One of these trials uses PCO as the initial condition, but this rarely gives the best solution. The solution is automatically rotated to the major axes (2D and 3D).

The algorithm implemented in PAST, which seems to work very well, is based on a new approach developed by Taguchi & Oono (in press).

The minimal spanning tree option is based on the selected similarity or distance index in the original space.

Shepard plot: This plot of obtained versus observed (target) ranks indicates the quality of the result. Ideally, all points should be placed on a straight ascending line ($x=y$). The R^2 values are the coefficients of determination between distances along each ordination axis and the original distances (perhaps not a very meaningful value, but is reported by other NMDS programs so is included for completeness).

Missing data is supported by pairwise deletion (not for the Raup-Crick, Rho and user-defined indices).

Correspondence analysis

Correspondence analysis (CA) is yet another ordination method, somewhat similar to PCA but for *counted data*. For comparing associations (columns) containing counts of taxa, or counted taxa (rows) across associations, CA is the more appropriate algorithm. Also, CA is more suitable if you expect that species have unimodal responses to the underlying parameters, that is they favour a certain range of the parameter, becoming rare for lower and higher values (this is in contrast to PCA, which assumes a linear response).

The CA routine finds the eigenvalues and eigenvectors of a matrix containing the Chi-squared distances between all data points. The eigenvalue, giving a measure of the similarity accounted for by the corresponding eigenvector, is given for each eigenvector. The percentages of similarity accounted for by these components are also given.

The 'View scatter' option allows you to see all your data points (rows) plotted in the coordinate system given by the CA. If you have colored (grouped) rows, the different groups will be shown using different symbols and colours.

In addition, the variables (columns, associations) can be plotted in the same coordinate system (Q mode), optionally including the column labels. If your data are 'well behaved', taxa typical for an association should plot in the vicinity of that association.

PAST presently uses a symmetric scaling ("Benzecri scaling").

If you have more than two columns in your data set, you can choose to view a scatter plot on the second and third axes.

Relay plot: This is a composite diagram with one plot per column. The plots are ordered according to CA column scores. Each data point is plotted with CA first-axis row scores on the vertical axis, and the original data point value (abundance) in the given column on the horizontal axis. This may be most useful when samples are in rows and taxa in columns. The relay plot will then show the taxa ordered according to their positions along the gradients, and for each taxon the corresponding plot should ideally show a unimodal peak, partly overlapping with the peak of the next taxon along the gradient (see Hennebert & Lees 1991 for an example from sedimentology).

Missing data is supported by column average substitution.

Reference

Hennebert, M. & A. Lees. 1991. Environmental gradients in carbonate sediments and rocks detected by correspondence analysis: examples from the Recent of Norway and the Dinantian of southwest England. *Sedimentology* 38:623-642.

Detrended correspondence analysis

The Detrended Correspondence (DCA) module uses the same algorithm as Decorana (Hill & Gauch 1980), with modifications according to Oxanen & Minchin (1997). It is specialized for use on 'ecological' data sets with abundance data; samples in rows, taxa in columns (vice versa prior to v. 1.79). When the 'Detrending' option is switched off, a basic Reciprocal Averaging will be carried out. The result should be similar to Correspondence Analysis (see above) plotted on the first and second axes.

Eigenvalues for the first three ordination axes are given as in CA, indicating their relative importance in explaining the spread in the data.

Detrending is a sort of normalization procedure in two steps. The first step involves an attempt to 'straighten out' points lying in an arch, which is a common occurrence. The second step involves 'spreading out' the points to avoid clustering of the points at the edges of the plot. Detrending may seem an arbitrary procedure, but can be a useful aid in interpretation.

Missing data is supported by column average substitution.

References

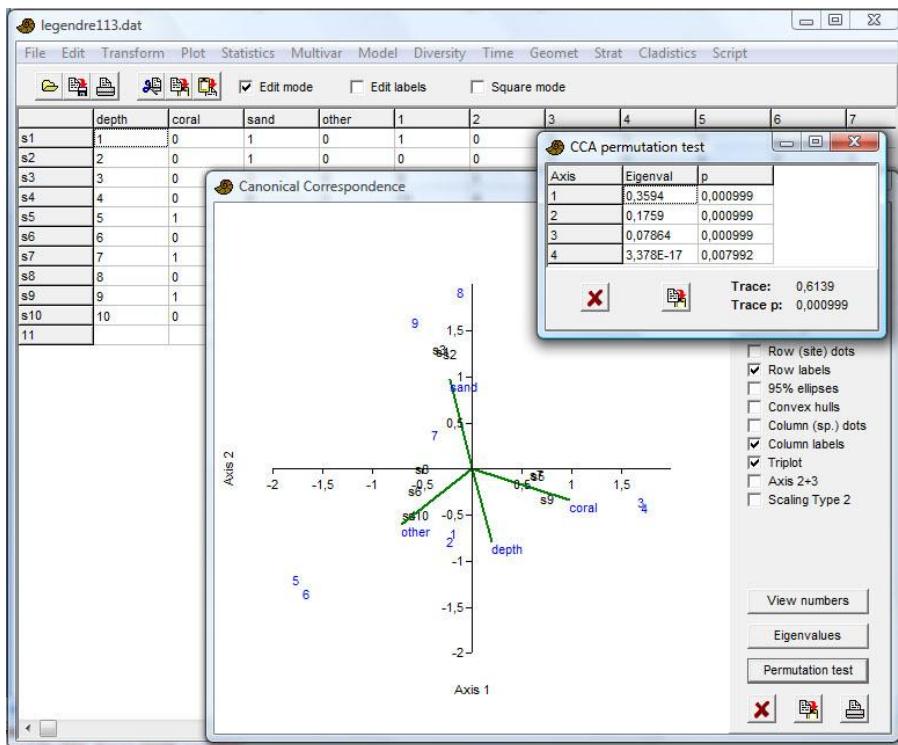
Hill, M.O. & H.G. Gauch Jr. 1980. Detrended Correspondence analysis: an improved ordination technique. *Vegetatio* 42:47-58.

Oxanen, J. & P.R. Minchin. 1997. Instability of ordination results under changes in input data order: explanations and remedies. *Journal of Vegetation Science* 8:447-454.

Canonical correspondence

Canonical Correspondence Analysis (Legendre & Legendre 1998) is correspondence analysis of a site/species matrix where each site has given values for one or more environmental variables (temperature, depth, grain size etc.). The ordination axes are linear combinations of the environmental variables. CCA is thus an example of direct gradient analysis, where the gradient in environmental variables is known *a priori* and the species abundances (or presence/absences) are considered to be a response to this gradient.

Each site should occupy one row in the spreadsheet. The environmental variables should enter in the first columns, followed by the abundance data (the program will ask for the number of environmental variables).



The implementation in PAST follows the eigenanalysis algorithm given in Legendre & Legendre (1998). The ordinations are given as site scores - fitted site scores are presently not available. Environmental variables are plotted as correlations with site scores. Both scalings (type 1 and 2) of Legendre & Legendre (1998) are available. Scaling 2 emphasizes relationships between species.

Reference

Legendre, P. & L. Legendre. 1998. Numerical Ecology, 2nd English ed. Elsevier, 853 pp.

CABFAC factor analysis

This module implements the classical Imbrie & Kipp (1971) method of factor analysis and environmental regression (CABFAC and REGRESS).

The program asks whether the first column contains environmental data. If not, a simple factor analysis with Varimax rotation will be computed on row-normalized data. If environmental data are included, the factors will be regressed onto the environmental variable using the second-order (parabolic) method of Imbrie & Kipp, with cross terms. PAST then reports the RMA regression of original environmental values against values reconstructed from the transfer function. You can also save the transfer function as a text file that can later be used for reconstruction of palaeoenvironment (see below). This file contains:

- Number of taxa
- Number of factors
- Factor scores for each taxon
- Number of regression coefficients
- Regression coefficients (second- and first-order terms, and intercept)

Reference

Imbrie & Kipp 1971.

Two-block PLS

Two-block Partial Least squares can be seen as an ordination method that can be compared with PCA, but with the objective of maximizing covariance between two sets of variates on the same rows (specimens, sites). For example, morphometric and environmental data collected on the same specimens can be ordinated in order to study covariation between the two.

The program will ask for the number of columns belonging to the first block. The remaining columns will be assigned to the second block. There are options for plotting PLS scores both within and across blocks, and PLS loadings.

The algorithm follows Rohlf & Corti (2000). Permutation tests and biplots are not yet implemented.

Missing data supported by column average substitution.

Reference

Rohlf, F.J. & M. Corti. 2000. Use of two-block partial least squares to study covariation in shape. *Systematic Biology* 49:740-753.

Cluster analysis

The hierarchical clustering routine produces a 'dendrogram' showing how data points (rows) can be clustered. For 'R' mode clustering, putting weight on groupings of taxa, taxa should go in rows. It is also possible to find groupings of variables or associations (Q mode), by entering taxa in columns. Switching between the two is done by transposing the matrix (in the Edit menu).

Three different algorithms are available:

- Unweighted pair-group average (UPGMA). Clusters are joined based on the average distance between all members in the two groups.
- Single linkage (nearest neighbour). Clusters are joined based on the smallest distance between the two groups.
- Ward's method. Clusters are joined such that increase in within-group variance is minimized,

One method is not necessarily better than the other, though single linkage is not recommended by some. It can be useful to compare the dendograms given by the different algorithms in order to informally assess the robustness of the groupings. If a grouping is changed when trying another algorithm, that grouping should perhaps not be trusted.

For Ward's method, a Euclidean distance measure is inherent to the algorithm. For UPGMA and single linkage, the distance matrix can be computed using 20 different indices, as described under the Statistics menu.

Missing data: The cluster analysis algorithm can handle missing data, coded as -1 or question mark (?). This is done using pairwise deletion, meaning that when distance is calculated between two points, any variables that are missing are ignored in the calculation. For Raup-Crick, missing values are treated as absence. Missing data are not supported for Ward's method, nor for the Rho similarity measure.

Two-way clustering: The two-way option allows simultaneous clustering in R-mode and Q-mode.

Stratigraphically constrained clustering: This option will allow only adjacent rows or groups of rows to be joined during the agglomerative clustering procedure. May produce strange-looking (but correct) dendograms.

Bootstrapping: If a number of bootstrap replicates is given (e.g. 100), the columns are subjected to resampling. The percentage of replicates where each node is still supported is given on the dendrogram.

Neighbour joining

Neighbour joining clustering (Saitou & Nei 1987) is an alternative method for hierarchical cluster analysis. The method was originally developed for phylogenetic analysis, but may be superior to UPGMA also for ecological data. In contrast with UPGMA, two branches from the same internal node do not need to have equal branch lengths. A phylogram (unrooted dendrogram with proportional branch lengths) is given.

Distance indices and bootstrapping are as for other cluster analysis (above).

Negative branch lengths are forced to zero, and transferred to the adjacent branch according to Kuhner & Felsenstein (1994).

The tree is by default rooted on the last branch added during tree construction (this is not midpoint rooting). Optionally, the tree can be rooted on the first row in the data matrix (outgroup).

References

Saitou, N. & M. Nei. 1987. The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Molecular Biology and Evolution* 4:406-425.

K-means clustering

K-means clustering (e.g. Bow 1984) is a non-hierarchical clustering method. The number of clusters to use is specified by the user, usually according to some hypothesis such as there being two sexes, four geographical regions or three species in the data set

The cluster assignments are initially random. In an iterative procedure, items are then moved to the cluster which has the closest cluster mean, and the cluster means are updated accordingly. This continues until items are no longer "jumping" to other clusters. The result of the clustering is to some extent dependent upon the initial, random ordering, and cluster assignments may therefore differ from run to run. This is not a bug, but normal behaviour in k-means clustering.

The cluster assignments may be copied and pasted back into the main spreadsheet, and corresponding colors (symbols) assigned to the items using the 'Numbers to colors' option in the Edit menu.

Missing data is supported by column average substitution.

Reference

Bow, S.-T. 1984. Pattern recognition. Marcel Dekker, New York.

Multivariate normality

Multivariate normality is assumed by a number of multivariate tests. PAST computes Mardia's multivariate skewness and kurtosis, with tests based on chi-squared (skewness) and normal (kurtosis) distributions. A powerful omnibus (overall) test due to Doornik & Hansen (1994) is also given. If at least one of these tests show departure from normality (small p value), the distribution is significantly non-normal. Sample size should be reasonably large (>50), although a small-sample correction is also attempted for the skewness test.

Discriminant/Hotelling

Given two sets of multivariate data, an axis is constructed which maximizes the difference between the sets. The two sets are then plotted along this axis using a histogram.

This module expects the rows in the two data sets to be grouped into two sets by coloring the rows, e.g. with black (dots) and red (crosses).

Equality of the means of the two groups is tested by a multivariate analogue to the *t* test, called *Hotelling's T-squared*, and a *p* value for this test is given. Normal distribution of the variables is required, and also that the number of cases is at least two more than the number of variables.

Number of constraints: For correct calculation of the Hotelling's *p* value, the number of dependent variables (constraints) must be specified. It should normally be left at 0, but for Procrustes fitted landmark data use 4 (for 2D) or 6 (for 3D).

Discriminant analysis can be used for visually confirming or rejecting the hypothesis that two species are morphologically distinct. Using a cutoff point at zero (the midpoint between the means of the discriminant scores of the two groups), a classification into two groups is shown in the "View numbers" option. The percentage of correctly classified items is also given.

Discriminant function: New specimens can be classified according to the discriminant function. Take the inner product between the measurements on the new specimen and the given discriminant function factors, and then subtract the given offset value.

Leave one out (cross-evaluation): An option is available for leaving out one row (specimen) at a time, re-computing the discriminant analysis with the remaining specimens, and classifying the left-out row accordingly (as given by the Score value).

Beware: The combination of discriminant analysis and Hotelling's T^2 test is sometimes misused. One should not be surprised to find a statistically significant difference between two samples which have been chosen with the objective of maximizing distance in the first place! The division into two groups should ideally be based on independent evidence.

See [Davis \(1986\)](#) for details.

Reference

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Paired Hotelling

The paired Hotelling's test expects two groups of multivariate data, marked with different colours. Rows within each group must be consecutive. The first row of the first group is paired with the first row of the second group, the second row is paired with the second, etc.

Missing data is supported by column average substitution.

Two-group permutation

This module expects the rows in the two data sets to be grouped into two sets by coloring the rows, e.g. with black (dots) and red (crosses).

Equality of the means of the two groups is tested using permutation with 2000 replicates (can be changed by the user), and the Mahalanobis squared distance measure. The permutation test is an alternative to Hotelling's test when the assumptions of multivariate normal distributions and equal covariance matrices do not hold.

Missing data is supported by column average substitution.

Box's M

This test is rather specialized, testing for the equivalence of the covariance matrices for two multivariate samples. You can use either two original multivariate samples from which the covariance matrices are automatically computed, or two specified variance-covariance matrices. In the latter case, you must also specify the sizes (number of individuals) of the two samples.

The Box's M statistic is given, together with a significance value based on a chi-square approximation. Note that this test is supposedly very sensitive. This means that a high p value will be a good, although informal, indicator of equality, while a highly significant result (low p value) may in practical terms be a somewhat too sensitive indicator of inequality.

MANOVA/CVA

One-way MANOVA (Multivariate ANalysis Of VAriance) is the multivariate version of the univariate ANOVA, testing whether several samples have the same mean. If you have only two samples, you would perhaps rather use the two-sample Hotelling's T^2 test.

Two statistics are provided: Wilk's lambda with it's associated Rao's F and the Pillai trace with it's approximated F. Wilk's lambda is probably more commonly used, but the Pillai trace may be more robust.

Number of constraints: For correct calculation of the p values, the number of dependent variables (constraints) must be specified. It should normally be left at 0, but for Procrustes fitted landmark data use 4 (for 2D) or 6 (for 3D).

Pairwise comparisons (post-hoc): If the MANOVA shows significant overall difference between groups, the analysis can proceed by pairwise comparisons. In PAST, the post-hoc analysis is quite simple, by pairwise Hotelling's tests. In the post-hoc table, groups are named according to the row label of the first item in the group. Hotelling's p values are given above the diagonal, while Bonferroni corrected values (multiplied by the number of pairwise comparisons) are given below the diagonal. This Bonferroni corrected test has very little power.

Canonical Variates Analysis

An option under MANOVA, CVA produces a scatter plot of specimens along the two first canonical axes, producing maximal and second to maximal separation between all groups (multigroup discriminant analysis). The axes are linear combinations of the original variables as in PCA, and eigenvalues indicate amount of variation explained by these axes.

Missing data is supported by column average substitution.

One-way ANOSIM

ANOSIM (ANalysis Of Similarities) is a non-parametric test of significant difference between two or more groups, based on any distance measure (Clarke 1993). The distances are converted to ranks. ANOSIM is normally used for ecological taxa-in-samples data, where groups of samples are to be compared.

Groups should be specified using row colors as usual.

In a rough analogy with ANOVA, the test is based on comparing distances between groups with distances within groups. Let rb be the mean rank of all distances between groups, and rw the mean rank of all distances within groups. The test statistic R is then defined as

$$R = (rb - rw) / (N(N-1)/4).$$

Large positive R (up to 1) signifies dissimilarity between groups. The significance is computed by permutation of group membership, with 10,000 replicates (can be changed by the user).

Pairwise ANOSIMs between all pairs of groups are provided as a post-hoc test. The Bonferroni corrected p values are very conservative.

Missing data is supported by pairwise deletion (not for the Raup-Crick, Rho and user-defined indices).

Reference

Clarke, K.R. 1993. Non-parametric multivariate analysis of changes in community structure. *Australian Journal of Ecology* 18:117-143.

Two-way ANOSIM

The two-way ANOSIM in PAST uses the crossed design (Clarke 1993). For more information see one-way ANOSIM, but note that groups (levels) are not coded with colors but with integer numbers in the first two columns.

Reference

Clarke, K.R. 1993. Non-parametric multivariate analysis of changes in community structure. *Australian Journal of Ecology* 18:117-143.

One-way NPMANOVA

NPMANOVA (Non-Parametric MANOVA) is a non-parametric test of significant difference between two or more groups, based on any distance measure (Anderson 2001). NPMANOVA is normally used for ecological taxa-in-samples data, where groups of samples are to be compared, but may also be used as a general non-parametric MANOVA

NPMANOVA calculates an *F* value in analogy with ANOVA. In fact, for univariate data sets and the Euclidean distance measure, NPMANOVA is equivalent to ANOVA and gives the same *F* value.

The significance is computed by permutation of group membership, with 10,000 replicates (can be changed by the user).

Reference

Anderson, M.J. 2001. A new method for non-parametric multivariate analysis of variance. *Austral Ecology* 26:32-46.

Mantel test

The Mantel test is a permutation test for correlation between two distance or similarity matrices. In PAST, these matrices can also be computed automatically from two sets of original data. The first matrix must be above the second matrix in the spreadsheet, and the rows be marked with two different colors. The two matrices must have the same number of rows. If they are distance or similarity matrices, they must also have the same number of columns.

SIMPER

SIMPER (Similarity Percentage) is a simple method for assessing which taxa are primarily responsible for an observed difference between groups of samples (Clarke 1993). The overall significance of the difference is often assessed by ANOSIM. The Bray-Curtis similarity measure is implicit to SIMPER.

If more than two groups are selected, you can either compare two groups (pairwise) by choosing from the lists of groups, or you can pool all samples to perform one overall multi-group SIMPER.

Reference

Clarke, K.R. 1993. Non-parametric multivariate analysis of changes in community structure. *Australian Journal of Ecology* 18:117-143.

Calibration from CABFAC

This module will reconstruct a (single) environmental parameter from taxa-in-samples abundance data. The program will also ask for a CABFAC transfer function file, as previously made by CABFAC factor analysis. The set of taxa (columns) must be identical in the spreadsheet and the transfer function file.

Calibration from optima

The first three rows can be generated from known (Recent) abundance and environmental data by the "Species packing" option in the Model menu. The third row (peak abundance) is not used, and the second row (tolerance) is used only when the "Equal tolerances" box is not ticked.

The algorithm is weighted averaging, optionally with tolerance weighting, according to ter Braak & van Dam (1989).

Reference

ter Braak, C.J.F & H. van Dam. 1989. Inferring pH from diatoms: a comparison of old and new calibration methods. *Hydrobiologia* 178:209-223.

Modern Analog Technique

The Modern Analog Technique works by finding modern samples with faunal associations close to modern ones.

Parameters to set:

- Distance measure: Several distance measures commonly used in MAT are available. "Squared chord" has become the standard choice in the literature.
- Weighting: When several modern analogs are linked to one downcore sample, their environmental values can be weighted equally, inversely proportional to faunal distance, or inversely proportional to ranked faunal distance.
- Distance threshold: Only modern analogs closer than this threshold are used. A default value is given, which is the tenth percentile of distances between all sample pairs in the modern data. The "Dissimilarity distribution" histogram may be useful when selecting this threshold.
- N analogs: This is the maximum number of modern analogs used for each downcore sample.
- Jump method (on/off): For each downcore sample, modern samples are sorted by ascending distance. When the distance increases by more than the selected percentage, the subsequent modern analogs are discarded.

Note that one or more of these options can be disabled by entering a large value. For example, a very large distance threshold will never apply, so the number of analogs is decided only by the "N analogs" value and optionally the jump method.

Cross validation

The scatter plot and R^2 value show the results of a leave-one-out (jackknifing) cross-validation within the modern data. The $y=x$ line is shown in red. This only partly reflects the "quality" of the method, as it gives little information about the accuracy of downcore estimation.

Model menu

Linear

If two columns are selected, they represent x and y values, respectively. If one column is selected, it represents y values, and x values are taken to be the sequence of positive integers (1,2,...). A straight line $y=ax+b$ is fitted to the data. There are two different algorithms available: Standard regression and Reduced Major Axis (the latter is selected by ticking the box). Standard regression keeps the x values fixed, and finds the line which minimizes the squared errors in the y values. Use this if your x values have very little error associated with them. Reduced Major Axis tries to minimize both the x and the y errors. RMA fitting and standard error estimation is according to Miller & Kahn (1962), *not* Davis (1986)!

Also, both x and y values can be log-transformed (base 10), in effect fitting your data to the 'allometric' function $y=10^b x^a$. An a value around 1 indicates that a straight-line ('isometric') fit may be more applicable.

The values for a and b , their errors, a Chi-square correlation value (not for RMA), Pearson's r correlation, and the probability that the columns are *not* correlated are given.

The calculation of standard errors for slope and intercept assumes normal distribution of residuals and independence between the variables and the variance of residuals. If these assumptions are strongly violated, it is preferable to use the bootstrapped 95 percent confidence intervals (2000 replicates). The number of random points selected for each replicate should normally be kept as N , but may be reduced for special applications.

The permutation test on correlation (R^2) uses 10,000 replicates.

In Standard regression (not RMA), a 95 percent "Working-Hotelling" confidence band for the fitted line (not for the data points!) is available.

Residuals

The Residuals window reports the distances from each data point to the regression line, in the x and y directions. Only the latter is of interest when using ordinary linear regression rather than RMA. The residuals can be copied back to the spreadsheet and inspected for normal distribution and independence between independent variable and residual variance (homoskedasticity).

Exponential functions

Your data can be fitted to an exponential function $y=e^b e^{ax}$ by first log-transforming just your y column (in the Transform menu) and then performing a straight-line fit.

References

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Miller, R.L. & Kahn, J.S. 1962. Statistical Analysis in the Geological Sciences. John Wiley & Sons, New York. 483 pp.

Linear, one independent, n dependent (multivariate regression)

When you have one independent variate and several dependent variates, you can fit each dependent variate separately to the independent variate using simple linear regression. This module makes the process more convenient by having a scroll button going through each dependent variate.

Two or more columns of measured data, with the independent in the first column and the dependents in consecutive columns.

In addition, an overall test of multivariate regression significance is provided (MANOVA with Wilks' lambda).

Linear, n independent, one dependent (multiple regression)

Two or more columns of measured data, with the dependent in the first column and the independents in consecutive columns.

The program will present the multiple correlation coefficient R and R^2 , together with the "adjusted" R^2 and an overall ANOVA-type significance test.

The coefficients (intercept, and slope for each independent variate) are presented with their estimated standard errors and t tests.

Polynomial regression

Two columns must be selected (x and y values). A polynomial of up to the fourth order is fitted to the data. The algorithm is based on a least-squares criterion and singular value decomposition (Press et al. 1992), with mean and variance standardization for improved numerical stability.

The chi-squared value is a measure of fitting error - larger values mean poorer fit. The Akaike Information Criterion has a penalty for the number of terms. The AIC should be as low as possible to maximize fit but avoid overfitting.

R^2 is the coefficient of determination, or proportion of variance explained by the model. Finally, a p value, based on an F test, gives the significance of the fit.

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

Sinusoidal regression

Two columns must be selected (x and y values). A sum of up to eight sinusoids with periods specified by the user, but with unknown amplitudes and phases, is fitted to the data. This can be useful for modeling periodicities in time series, such as annual growth cycles or climatic cycles, usually in combination with spectral analysis. The algorithm is based on a least-squares criterion and singular value decomposition (Press et al. 1992). By default, the periods are set to the range of the x values, and harmonics (1/2, 1/3, 1/4, 1/5, 1/6, 1/7 and 1/8 of the fundamental period). These values can be changed, and need not be in harmonic proportion.

The chi-squared value is a measure of fitting error - larger values mean poorer fit. The Akaike Information Criterion has a penalty for the number of sinusoids (the equation used assumes that the periods are estimated from the data). The AIC should be as low as possible to maximize fit but avoid overfitting.

R^2 is the coefficient of determination, or proportion of variance explained by the model. Finally, a p value, based on an F test, gives the significance of the fit.

A "search" function for each sinusoid will optimize the frequency of that sinusoid (over the full meaningful range from one period to the Nyquist frequency), holding all other selected sinusoid frequencies constant. The algorithm is slow but very robust and almost guaranteed to find the global optimum.

For a "blind" spectral analysis, finding all parameters of an optimal number of sinusoids, follow this procedure: Start with only the first sinusoid selected. Click "search" to optimize period, amplitude and phase. This will find the strongest sinusoid in the data. Note the AIC. Add (select) the second sinusoid, and click its search button to optimize all parameters of both sinusoids except the period of the first sinusoid. This will find the second strongest sinusoid. Continue until the AIC no longer decreases.

It is not meaningful to specify periodicities that are smaller than two times the typical spacing of data points.

Each sinusoid is given by $y=a\cos(2\pi(x-x_0)/T - p)$, where a is the amplitude, T is the period and p is the phase. x_0 is the first (smallest) x value.

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

Logistic

Attempts to fit the data to the logistic equation $y=a/(1+be^{-cx})$, where x starts at zero (that is, your data set is shifted so that your leftmost data point is moved to $x=0$). The algorithm is a little complicated. The value of a is first estimated to be the maximal value of y . The values of b and c are then estimated using a straight-line fit to a linearized model.

Though acceptable, this estimate can optionally be improved by using the estimated values as an initial guess for a Levenberg-Marquardt nonlinear optimization (Press et al. 1992). This procedure can sometimes improve the fit, but due to the numerical instability of the logistic model it can fail with an error message.

The logistic equation can model growth with saturation, and was used by Sepkoski (1984) to describe the proposed stabilization of marine diversity in the late Palaeozoic.

The 95% confidence intervals are based on 2000 bootstrap replicates, not using the Levenberg-Marquardt optimization step.

Von Bertalanffy

An option in the 'Logistic fit' window. Uses the same algorithm as above, but fits to the von Bertalanffy equation $y=a(1-be^{-cx})$. This equation is used for modelling growth of multi-celled animals (in units of length or width, not volume).

More information about the logistic and the von Bertalanffy growth models can be found in Brown & Rothery (1993).

References

Brown, D. & P. Rothery. 1993. Models in biology: mathematics, statistics and computing. John Wiley & Sons.

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

Sepkoski, J.J. 1984. A kinetic model of Phanerozoic taxonomic diversity. *Paleobiology* 10:246-267.

Smoothing spline

Two columns must be selected (x and y values). The data are fitted to a smoothing spline, which is a sequence of third-order polynomials, continuous up to the second derivative. A typical application of this is the construction of a smooth curve going through a noisy data set. The algorithm is due to de Boor (2001).

An optional third column specifies standard deviations on the data points. These are used for weighting the data. If unspecified, these values are all set to 10% of the standard deviation of the Y values.

A smoothing value is set by the user. This is a normalized version of the smoothing factor of de Boor (default 1). Larger values give smoother curves. A value of 0 will start a spline segment at every point. Clicking "Optimize smoothing" will calculate an "optimal" smoothing factor by a crossvalidation (jackknife) procedure.

"View given points" gives a table of the given data points X, Y and stdev(Y), the corresponding Y values on the spline curve (ys) and the residuals. The chi-squared test for each point may be used to identify outliers. The final column suggests an stdev(Y) value to use if forcing the p value to 0.5.

An optional fourth column (if used then the third column must also be filled with stdev values) may not contain the same number of values as the previous columns. It contains extra X values to be used for interpolation between the data points.

Note that sharp jumps in your data can give rise to oscillations in the curve, and that you can also get large excursions in regions with few data points.

Multiple data points at the same X value are handled by collapsing them to a single point by weighted averaging and calculation of a combined standard deviation.

Reference

de Boor, Carl. 2001. A practical guide to splines. Springer.

Abundance models

This module can be used for plotting taxon abundances in descending rank order on a linear or logarithmic (Whittaker plot) scale, or number of species in abundance octave classes (as shown when fitting to log-normal distribution). It can also fit the data to one of four different standard abundance models:

- Geometric, where the 2nd most abundant species should have a taxon count of $k < 1$ times the most abundant, the 3rd most abundant a taxon count of k times the 2nd most abundant etc. for a constant k . This will give a straight descending line in the Whittaker plot. Fitting is by simple linear regression of the log abundances.
- Log-series, with two parameters α and x . The fitting algorithm is from Krebs (1989).
- Broken stick. There are no free parameters to be fitted in this model.
- Log-normal. The fitting algorithm is from Krebs (1989). The logarithm (base 10) of the fitted mean and variance are given. The *octaves* refer to power-of-2 abundance classes:

Octave	Abundance
1	1
2	2-3
3	4-7
4	8-15
5	16-31
6	32-63
7	64-127
...	...

A significance value based on chi-squared is given for each of these models, but the power of the test is not the same for the four models and the significance values should therefore not be compared. It is important, as always, to remember that a high p value can not be taken to imply a good fit. A low value does however imply a bad fit.

Reference

Krebs, C.J. 1989. Ecological Methodology. Harper & Row, New York.

Species packing (Gaussian)

This module fits Gaussian response models to species abundances along a gradient, for one or more species. The fitted parameters are optimum (average), tolerance (standard deviation) and maximum.

One column of environmental measurements in samples (e.g. temperature), and one or more columns of abundance data (taxa in columns).

Two algorithms are available: Least-squares Gaussian fitting by internal transformation to a parabolic form, and weighted averaging according to ter Braak & van Dam (1989).

Reference

ter Braak, C.J.F & H. van Dam. 1989. Inferring pH from diatoms: a comparison of old and new calibration methods. *Hydrobiologia* 178:209-223.

Diversity menu

Diversity indices

These statistics apply to association data, where number of individuals are tabulated in rows (taxa) and possibly several columns (associations). The available statistics are as follows, for each association:

Number of taxa (S)

Total number of individuals (n)

Dominance=1-Simpson index. Ranges from 0 (all taxa are equally present) to 1 (one taxon dominates the community completely).

$D=\sum((n_i/n)^2)$ where n_i is number of individuals of taxon i .

Simpson index=1-dominance. Measures 'evenness' of the community from 0 to 1. Note the confusion in the literature: Dominance and Simpson indices are often interchanged!

Shannon index (entropy). A diversity index, taking into account the number of individuals as well as number of taxa. Varies from 0 for communities with only a single taxon to high values for communities with many taxa, each with few individuals.

$H=-\sum((n_i/n)\ln(n_i/n))$

Buzas and Gibson's evenness = e^H/S

Menhinick's richness index - the ratio of the number of taxa to the square root of sample size.

Margalef's richness index: $(S-1)/\ln(n)$, where S is the number of taxa, and n is the number of individuals.

Equitability. Shannon diversity divided by the logarithm of number of taxa. This measures the evenness with which individuals are divided among the taxa present.

Fisher's alpha - a diversity index, defined implicitly by the formula $S=a*\ln(1+n/a)$ where S is number of taxa, n is number of individuals and a is the Fisher's alpha.

Berger-Parker dominance: simply the number of individuals in the dominant taxon relative to n .

Many of these indices are explained in Harper (1999).

Approximate confidence intervals for all these indices can be computed with a bootstrap procedure. 1000 random samples are produced (200 prior to version 0.87b), each with the same total number of individuals as in each original sample. The random samples are taken from the total, pooled data set (all columns). For each individual in the random sample, the taxon is chosen with probabilities according to the original, pooled abundances. A 95 percent confidence interval is then calculated. Note that the diversity in the replicates will often be less than, and never larger than, the pooled diversity in the total data set.

Since these confidence intervals are all computed with respect to the pooled data set, they do not represent confidence intervals for the individual samples. They are mainly useful for identifying samples where the given diversity index falls outside the confidence interval. Bootstrapped comparison of diversity indices in two samples is provided in the Compare diversities module.

Reference

Harper, D.A.T. (ed.). 1999. Numerical Palaeobiology. John Wiley & Sons.

Quadrat richness

Two or more columns, each containing presence/absence (1/0) of different taxa down the rows.

Four non-parametric species richness estimators are included in PAST: Chao 2, first- and second-order jackknife, and bootstrap. All of these require presence-absence data in two or more sampled quadrats of equal size. Colwell & Coddington (1994) reviewed these estimators, and found that the Chao2 and the second-order jackknife performed best.

Reference

Colwell, R.K. & J.A. Coddington. 1994. Estimating terrestrial biodiversity through extrapolation. Philosophical Transactions of the Royal Society (Series B) 345:101-118.

Beta diversity

Two or more rows (samples) of presence-absence (0/1) data, with taxa in columns.

The beta diversity module in Past can be used for any number of samples (not limited to only two samples). The eight measures available are described in Koleff et al. (2003), and the table below refers to their notation:

Past	Koleff et al.
Whittaker	b_w
Harrison	b_{-1}
Cody	b_c
Routledge	b_l
Wilson-Shmida	b_t
Mourelle	b_{me}
Harrison 2	b_{-2}
Williams	b_{-3}

Reference

Koleff, P., K.J. Gaston & J.J. Lennon. 2003. Measuring beta diversity for presence-absence data. *Journal of Animal Ecology* 72:367-382.

Taxonomic distinctness

One or more columns, each containing counts of individuals of different taxa down the rows. In addition, the leftmost row(s) must contain names of genera/families etc. (see below).

Taxonomic diversity and taxonomic distinctness as defined by Clarke & Warwick (1998), including confidence intervals computed from 200 random replicates taken from the pooled data set (all columns). Note that the "global list" of Clarke & Warwick is not entered directly, but is calculated internally by pooling (summing) the given samples.

These indices depend on taxonomic information also above the species level, which has to be entered for each species as follows. Species names go in the name column (leftmost, fixed column), genus names in column 1, family in column 2 etc. Species counts follow in the columns thereafter. The program will ask for the number of columns containing taxonomic information above the species level.

For presence-absence data, taxonomic diversity and distinctness will be valid but equal to each other.

Reference

Clarke, K.R. & Warwick, R.M. 1998. A taxonomic distinctness index and its statistical properties. *Journal of Applied Ecology* 35:523-531.

Individual rarefaction

Comparing taxonomical diversity in samples of different sizes.

When comparing samples: Samples should be taxonomically similar, obtained using standardised sampling and taken from similar 'habitat'.

One or more columns of counts of individuals of different taxa (each column must have the same number of values)

Given one or more columns of abundance data for a number of taxa, this module estimates how many taxa you would expect to find in a sample with a smaller total number of individuals. With this method, you can compare the number of taxa in samples of different size. Using rarefaction analysis on your *largest* sample, you can read out the number of expected taxa for any smaller sample size (including that of the *smallest* sample). The algorithm is from Krebs (1989). An example application in paleontology can be found in Adrain et al. (2000).

Let N be the total number of individuals in the sample, s the total number of species, and N_i the number of individuals of species number i . The expected number of species $E(S_n)$ in a sample of size n and the variance $V(S_n)$ are then given by

$$E(S_n) = \sum_{i=1}^s \left[1 - \frac{\binom{N - N_i}{n}}{\binom{N}{n}} \right]$$
$$V(S_n) = \sum_{i=1}^s \left[\frac{\binom{N - N_i}{n}}{\binom{N}{n}} \left(1 - \frac{\binom{N - N_i}{n}}{\binom{N}{n}} \right) \right] + 2 \sum_{j=2}^s \sum_{i=1}^{j-1} \left[\frac{\binom{N - N_i - N_j}{n}}{\binom{N}{n}} - \frac{\binom{N - N_i}{n} \binom{N - N_j}{n}}{\binom{N}{n} \binom{N}{n}} \right]$$

Standard errors (square roots of variances) are given by the program. In the graphical plot, these standard errors are converted to 95 percent confidence intervals.

References

Adrain, J.M., S.R. Westrop & D.E. Chatterton. 2000. Silurian trilobite alpha diversity and the end-Ordovician mass extinction. *Paleobiology* 26:625-646.

Krebs, C.J. 1989. Ecological Methodology. Harper & Row, New York.

Sample rarefaction (Mao tau)

A matrix of presence-absence data (abundances treated as presences), with taxa in rows and samples in columns.

Sample-based rarefaction (also known as the species accumulation curve) is applicable when a number of samples are available, from which species richness is to be estimated as a function of number of samples. PAST implements the analytical solution known as "Mao tau", with standard deviation. In the graphical plot, the standard errors are converted to 95 percent confidence intervals.

See Colwell et al. (2004) for details.

Reference

Colwell, R.K., C.X. Mao & J. Chang. 2004. Interpolating, extrapolating, and comparing incidence-based species accumulation curves. *Ecology* 85:2717-2727.

Compare diversities

Equal sampling conditions

Two columns of abundance data with taxa down the rows

This module computes a number of diversity indices for two samples, and then compares the diversities using two different randomisation procedures as follows.

Bootstrapping

The two samples A and B are pooled. 1000 random pairs of samples (A_i, B_i) are then taken from this pool, with the same numbers of individuals as in the original two samples. For each replicate pair, the diversity indices $\text{div}(A_i)$ and $\text{div}(B_i)$ are computed. The number of times $|\text{div}(A_i) - \text{div}(B_i)|$ exceeds or equals $|\text{div}(A) - \text{div}(B)|$ indicates the probability that the observed difference could have occurred by random sampling from one parent population as estimated by the pooled sample.

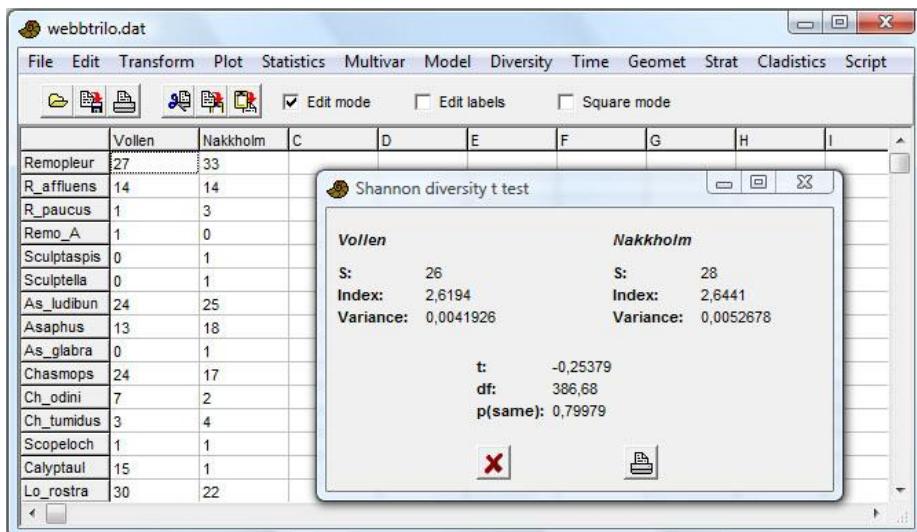
A small probability value $p(\text{same})$ then indicates a significant difference in diversity index between the two samples.

Permutation

1000 random matrices with two columns (samples) are generated, each with the same row and column totals as in the original data matrix. The p value is computed as for the bootstrap test.

Diversity *t* test

Comparison of the Shannon diversities (entropies) in two samples, using a *t* test described by e.g. Poole (1974) and Magurran (1988). This is an alternative to the randomization test available in the *Compare diversities* module. Two columns of abundance data with taxa down the rows are required.



Note that the Shannon indices here include a bias correction term (Poole 1974), and may diverge slightly from the uncorrected estimates calculated elsewhere in PAST, at least for small samples.

With p_i the proportion (0-1) of taxon i , S the number of taxa and N the number of individuals, the estimator of the index is computed as

$$H' = -\sum_{i=1}^S p_i \ln p_i - \frac{S-1}{2N} \text{ (note that the second term is incorrect in Magurran 1988).}$$

The variance of the estimator is

$$\text{Var } H' = \frac{\sum p_i (\ln p_i)^2 - \sum (p_i \ln p_i)^2}{N} - \frac{S-1}{2N^2}.$$

The *t* test statistic is given by

$$t = \frac{H'_1 - H'_2}{\sqrt{\text{Var } H'_1 + \text{Var } H'_2}}.$$

The degrees of freedom for the *t* test is

$$df = \frac{\frac{(\text{Var } H'_1 + \text{Var } H'_2)^2}{N_1}}{\frac{(\text{Var } H'_1)^2}{N_1} + \frac{(\text{Var } H'_2)^2}{N_2}}.$$

References

Magurran, A. 1988. Ecological Diversity and Its Measurement. Princeton University Press.

Poole, R.W. 1974. An introduction to quantitative ecology. McGraw-Hill, New York.

Diversity profiles

Two columns of abundance data with taxa down the rows

Equal sampling conditions

Comparing diversities in two samples of abundance data

The validity of comparing diversities in two samples can be criticized because of arbitrary choice of diversity index. One sample may for example contain a larger number of taxa, while the other has a larger Shannon index. It may therefore be a good idea to try a number of diversity indices in order to make sure that the diversity ordering is robust. A formal way of doing this is to define a family of diversity indices, dependent upon a single continuous parameter (Tothmeresz 1995).

PAST uses the exponential of the so-called Renyi index, which depends upon a parameter alpha. For alpha=0, this function gives the total species number. alpha=1 gives an index proportional to the Shannon index, while alpha=2 gives an index which behaves like the Simpson index.

The program plots two such diversity profiles together. If the profiles cross, the diversities are non-comparable.

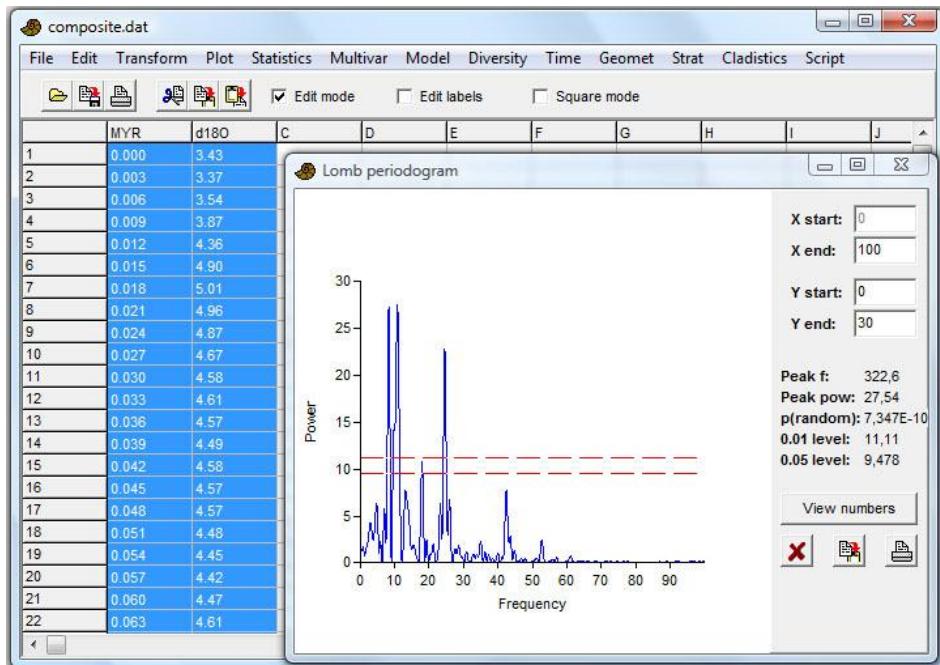
Reference

Tothmeresz, B. 1995. Comparison of different methods for diversity ordering. *Journal of Vegetation Science* 6:283-290.

Time series menu

Spectral analysis

Since palaeontological data are often unevenly sampled, Fourier-based methods can be difficult to use. PAST therefore includes the Lomb periodogram algorithm for unevenly sampled data (Press et al. 1992), with time values given in the first column and dependent values in the second column. If only one column is selected, an even spacing of one unit between data points is assumed. The Lomb periodogram should then give similar results as the FFT.



The frequency axis is in units of $1/(x \text{ unit})$. If for example, your x values are given in millions of years, a frequency of 0.1 corresponds to a period of 10 million years. The power axis is in units proportional to the square of the amplitudes of the sinusoids present in the data.

Also note that the frequency axis extends to very high values. If your data are evenly sampled, the upper half of the spectrum is a mirror image of the lower half, and is of little use. If some of your regions are closely sampled, the algorithm may be able to find useful information even above the half-point (Nyquist frequency).

The highest peak in the spectrum is presented with its frequency and power value, together with a probability that the peak could occur from random data. The 0.01 and 0.05 significance levels ('white noise lines') are shown as red dashed lines.

The example above shows a spectral analysis of a foram oxygen isotope record from 1 Ma to Recent, with an even spacing of 0.003 Ma (3 ka). There are periodicities at frequencies of about 9 (split peak), 25 and 43 Ma⁻¹, corresponding to periods of 111 ka, 40 ka and 23 ka – clearly orbital forcing.

Detrending

The data are not automatically detrended. It is recommended to detrend the data prior to spectral analysis, by subtraction of the linear regression line (*Remove trend* in the *Transform* menu). This will usually improve resolution at low frequencies.

Reference

Press, W.H., S.A. Teukolsky, W.T. Vetterling & B.P. Flannery. 1992. Numerical Recipes in C. Cambridge University Press.

Autocorrelation

Autocorrelation (Davis 1986) is carried out on separate column(s) of *evenly sampled* temporal/stratigraphic data. Lag times τ up to $n/2$, where n is the number of values in the vector, are shown along the x axis (positive lag times only - the autocorrelation function is symmetrical around zero). A predominantly zero autocorrelation signifies random data - periodicities turn up as peaks.

The "95 percent confidence interval" option will draw lines at

$$\pm 1.76 \sqrt{\frac{1}{n - \tau + 3}}$$

after Davis (1986). This is the confidence interval for random, independent points (white noise). There are two issues: White noise is an unrealistic null model, and the confidence interval is only strictly valid at each *individual* lag (multiple testing problem).

Missing data supported.

Reference

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Cross-correlation

Cross-correlation (Davis 1986) is carried out on two column(s) of *evenly sampled* temporal/stratigraphic data. The x axis shows the displacement of the second column with respect to the first, the y axis the correlation between the two time series for a given displacement. The "p values" option will draw the significance of the correlation, after Davis (1986).

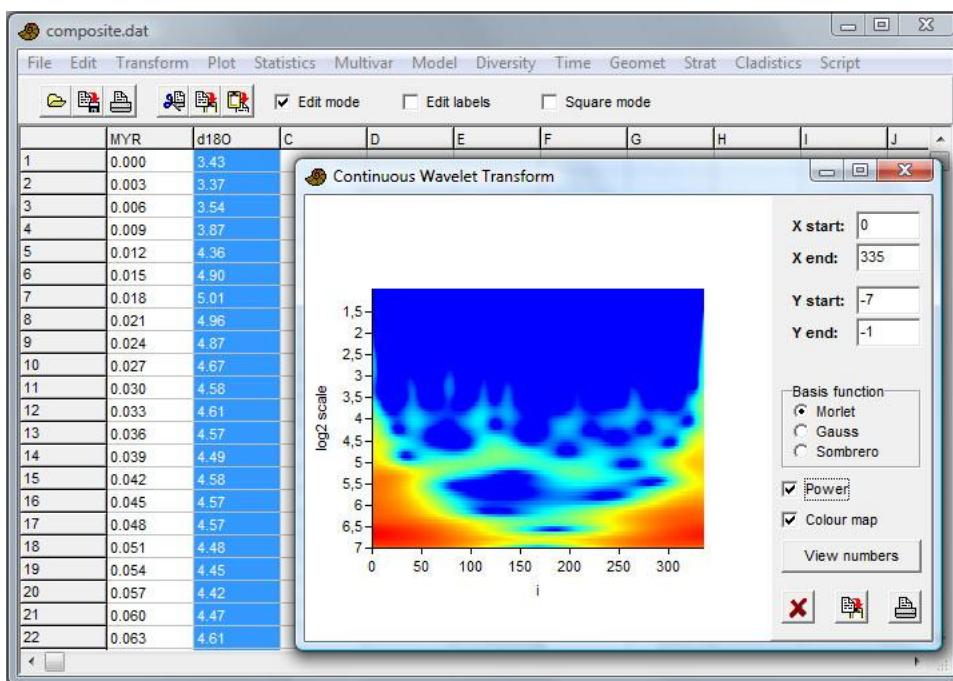
Reference

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Wavelet transform

Inspection of time series at different scales. Even spacing of data points. One column of counted or measured data

The continuous wavelet transform (CWT) is an analysis method where a data set can be inspected at small, intermediate and large scales simultaneously. It can be useful for detecting periodicities at different wavelengths, self-similarity and other features. The vertical axis in the plot is a logarithmic size scale (base 2), with the signal observed at a scale of only two consecutive data points at the top, and at a scale of one fourth of the whole sequence at the bottom. One unit on this axis corresponds to a doubling of the size scale. The top of the figure thus represents a detailed, fine-grained view, while the bottom represents a smoothed overview of longer trends. Signal power (or more correctly squared correlation strength with the scaled mother wavelet) is shown with a grayscale or in colour.

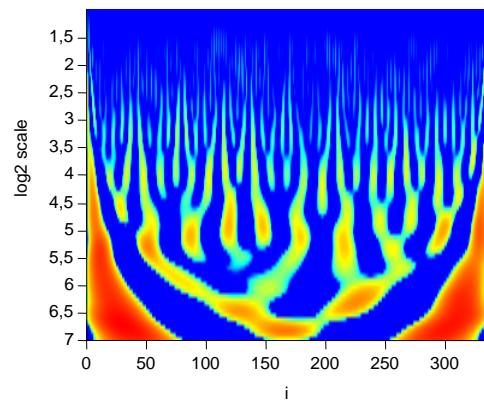


The shape of the mother wavelet can be set to Morlet, Gauss or Sombrero. The Morlet wavelet usually performs best.

The example above is based on a foram oxygen isotope record from 1 Ma to Recent, with an even spacing of 0.003 Ma (3 ka). A band can be seen at a scale of about $2^5=32$ samples, or about 100 ka. A weaker band around $2^4=16$ samples corresponds to a scale of about 48 ka. These are orbital periodicities. In contrast with the “bulk” spectral analysis, the scalogram makes visible changes in strength and frequency over time. Note the high power for large scales at the beginning and end of the time series – this is an artefact due to proximity to the boundary (so-called “cone of influence”).

The algorithm is based on fast convolution of the signal with the wavelet at different scales, using the FFT.

If the “Power” option is deselected, the program will show only the real part of the scalogram (not squared). This shows the signal in the time domain, filtered at different scales:



In the 'View numbers' window, sample number (position) is shown in the first row, and the scale values (\log_2) in the first column.

The wavelet transform was used by Prokoph et al. (2000) for illustrating cycles in diversity curves for planktic foraminifera.

Morlet basis: Complex exponential with a Gaussian envelope:

Reference

Prokoph, A., A.D. Fowler & R.T. Patterson. 2000. Evidence for periodicity and nonlinearity in a high-resolution fossil record of long-term evolution. *Geology* 28:867-870.

Walsh transform

Spectral analysis (finding periodicities) of binary or ordinal data. Even spacing of data points. One column of binary (0/1) or ordinal (integer) data

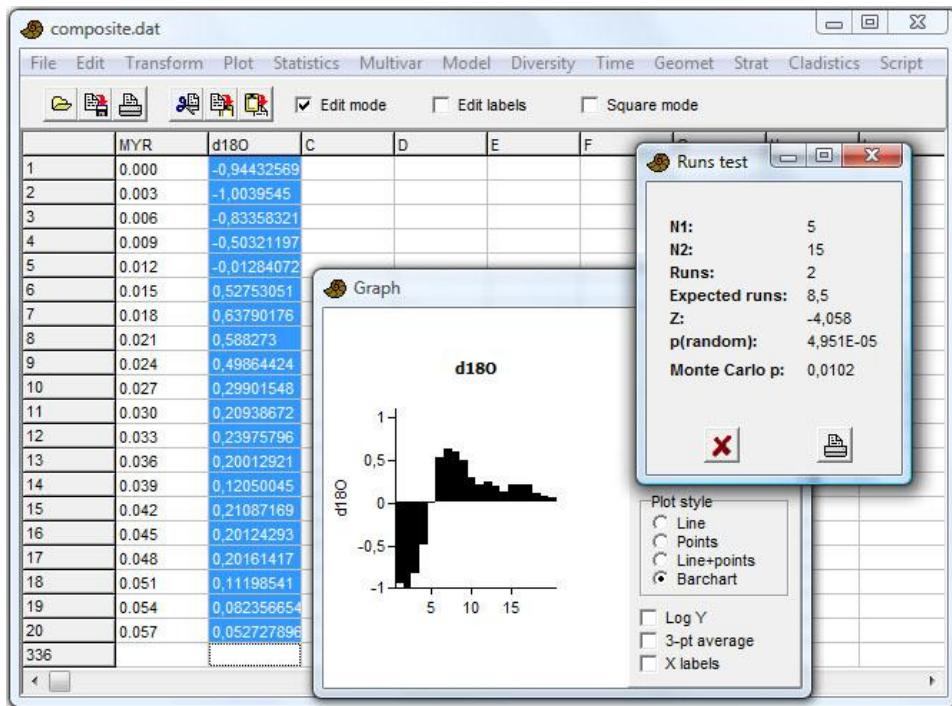
The normal methods for spectral analysis are perhaps not optimal for binary data, because they decompose the time series into sinusoids rather than "square waves". The Walsh transform may then be a better choice, using basis functions that flip between -1 and +1. These basis functions have different "frequencies" (number of transitions divided by two), known as *sequencies*. In PAST, each pair of even ("cal") and odd ("sal") basis functions (one pair for each integer-valued sequency) is combined into a power value using $\text{cal}^2 + \text{sal}^2$, producing a "power spectrum" that is comparable to the Lomb periodogram.

Note that the Walsh transform is slightly "exotic" compared with the Fourier transform, and its interpretation must be done cautiously. For example, the effects of the duty cycle (percentage of ones versus zeros) are somewhat difficult to understand.

In PAST, the data values are pre-processed by multiplying with two and subtracting one, bringing 0/1 binary values into the -1/+1 range optimal for the Walsh transform.

Runs test

The runs test is a non-parametric test for randomness in a sequence of values such as a time series. Non-randomness may include such effects as autocorrelation, trend and periodicity. The module requires one column of data, which are internally converted to 0 ($x \leq 0$) or 1 ($x > 0$).



The test is based on a dichotomy between two values ($x \leq 0$ or $x > 0$). It counts the number of runs (groups of consecutive equal values) and compares this to a theoretical value. The runs test can therefore be used directly for sequences of binary data. Continuous data can be converted to an appropriate form by subtracting the mean (*Transform* menu), or taking the difference from one value to the next (use "x-u" in the *Evaluate Expression* function).

With n the total number of data points, n_1 the number of points ≤ 0 and n_2 the number of points > 0 , the expected number of runs in a random sequence is

$$E(R) = \frac{n + 2n_1 n_2}{n}.$$

The variance is

$$Var(R) = \frac{2n_1 n_2 (2n_1 n_2 - n)}{n^2 (n-1)}.$$

With the observed number of runs R , a z statistic can be written as

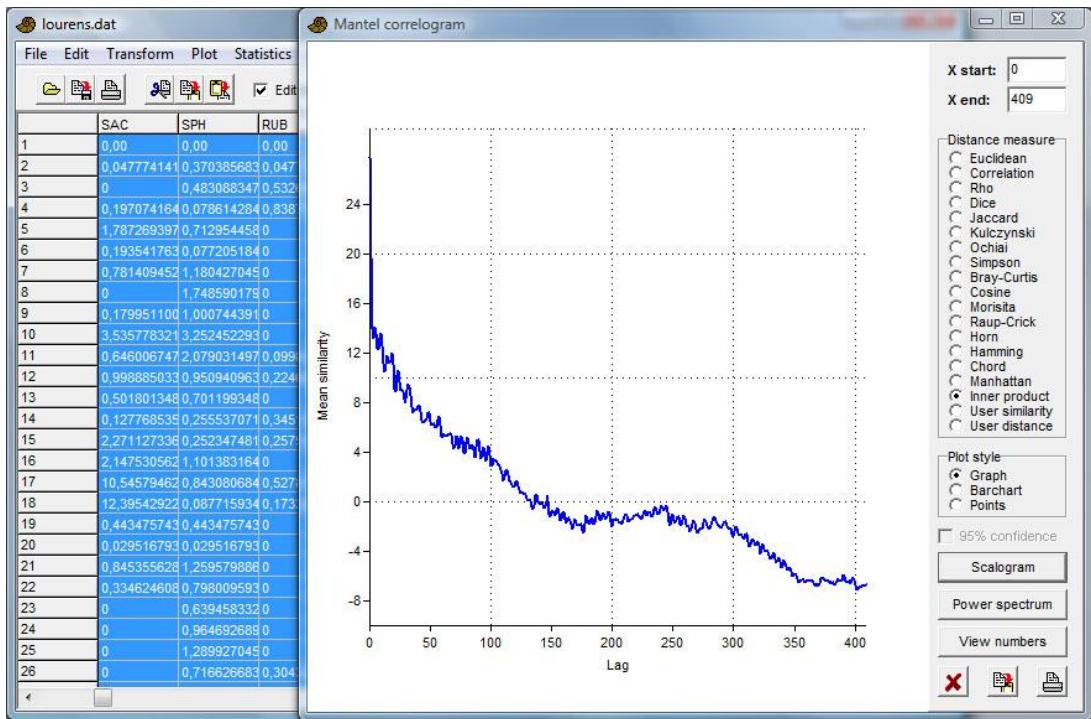
$$z = \frac{R - E(R)}{\sqrt{Var(R)}}.$$

The resulting two-tailed p value is not accurate for $n < 20$. A Monte Carlo procedure is therefore also included, based on 10,000 random replicates using the observed n , n_1 and n_2 .

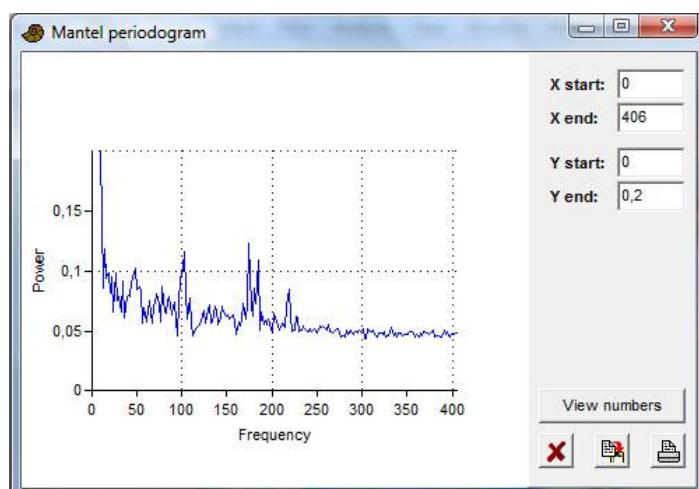
Mantel correlogram (and periodogram)

Several rows of multivariate data, one row for each sample. Samples are assumed to be evenly spaced in time.

The Mantel correlogram (e.g. Legendre & Legendre 1998) is a multivariate extension to autocorrelation, based on any similarity or distance measure. The Mantel correlogram in PAST shows the average similarity between the time series and a time lagged copy, for different lags.



The Mantel periodogram is a power spectrum of the multivariate time series, computed from the Mantel correlogram (Hammer 2007).



The Mantel scalogram is an experimental plotting of similarities between all pairs of points along the time series. The apex of the triangle is the similarity between the first and last point. The base of the triangle shows similarities between pairs of consecutive points.

References

Hammer, Ø. 2007. Spectral analysis of a Plio-Pleistocene multispecies time series using the Mantel periodogram. *Palaeogeography, Palaeoclimatology, Palaeoecology* 243:373-377.

Legendre, P. & L. Legendre. 1998. Numerical Ecology, 2nd English ed. Elsevier, 853 pp.

ARMA (and intervention analysis)

Analysis and removal of serial correlations in time series, and analysis of the impact of an external disturbance ("intervention") at a particular point in time. Stationary time series, except for a single intervention. One column of equally spaced data.

This powerful but somewhat complicated module implements maximum-likelihood ARMA analysis, and a minimal version of Box-Jenkins intervention analysis (e.g. for investigating whether a climate change had an impact on biodiversity).

By default, a simple ARMA analysis without interventions is computed. The user selects the number of AR (autoregressive) and MA (moving-average) terms to include in the ARMA difference equation. The log-likelihood and Akaike information criterion are given. Select the numbers of terms that minimize the Akaike criterion, but be aware that AR terms are more "powerful" than MA terms. Two AR terms can model a periodicity, for example.

The main aim of ARMA analysis is to remove serial correlations, which otherwise cause problems for model fitting and statistics. The residual should be inspected for signs of autocorrelation, e.g. by copying the residual from the numerical output window back to the spreadsheet and using the autocorrelation module. Note that for many paleontological data sets with sparse data and confounding effects, proper ARMA analysis (and therefore intervention analysis) will be impossible.

The program is based on the likelihood algorithm of Melard 1983, combined with nonlinear multivariate optimization using simplex search.

Intervention analysis proceeds as follows. First, carry out ARMA analysis on only the samples preceding the intervention, by typing the last pre-intervention sample number in the "last samp" box. It is also possible to run the ARMA analysis only on the samples following the intervention, by typing the first post-intervention sample in the "first samp" box, but this is not recommended because of the post-intervention disturbance. Also tick the "Intervention" box to see the optimized intervention model.

The analysis follows Box and Tiao (1976) in assuming an "indicator function" $u(i)$ that is either a unit step or a unit pulse, as selected by the user. The indicator function is transformed by an AR(1) process with a parameter δ , and then scaled by a magnitude (note that the magnitude given by PAST is the coefficient on the transformed indicator function: first do $y(i)=\delta*y(i-1)+u(i)$, then scale y by the magnitude). The algorithm is based on ARMA transformation of the complete sequence, then a corresponding ARMA transformation of y , and finally linear regression to find the magnitude. The parameter δ is optimized by exhaustive search over $[0,1]$.

For small impacts in noisy data, δ may end up on a sub-optimum. Try both the step and pulse options, and see what gives smallest standard error on the magnitude. Also, inspect the "delta optimization" data, where standard error of the estimate is plotted as a function of δ , to see if the optimized value may be unstable.

The Box-Jenkins model can model changes that are abrupt and permanent (step function with $\delta=0$, or pulse with $\delta=1$), abrupt and non-permanent (pulse with $\delta<1$), or gradual and permanent (step with $\delta<0$).

Be careful with the standard error on the magnitude - it will often be underestimated, especially if the ARMA model does not fit well. For this reason, a p value is deliberately not computed (Murtaugh 2002).

References

Box & Tiao 1976.

Melard 1983.

Murtaugh 2002.

Insolation (solar forcing) model

This module computes solar insolation at any latitude and any time from 100 Ma to the Recent (the results are less accurate before 50 Ma). The calculation can be done for a "true" orbital longitude, "mean" orbital longitude (corresponding to a certain date in the year), averaged over a certain month in each year, or integrated over a whole year.

The implementation in PAST is ported from the code by Laskar et al. (2004), by courtesy of these authors. Please reference Laskar et al. (2004) in any publications.

It is necessary to specify a data file containing orbital parameters. Download the file INSOLN.LA2004.BTL.100.ASC from <http://www.imcce.fr/Equipes/ASD/insola/earth/La2004> and put in anywhere on your computer. The first time you run the calculation, PAST will ask for the position of the file.

The amount of data can become excessive for long time spans and short step sizes!

Reference

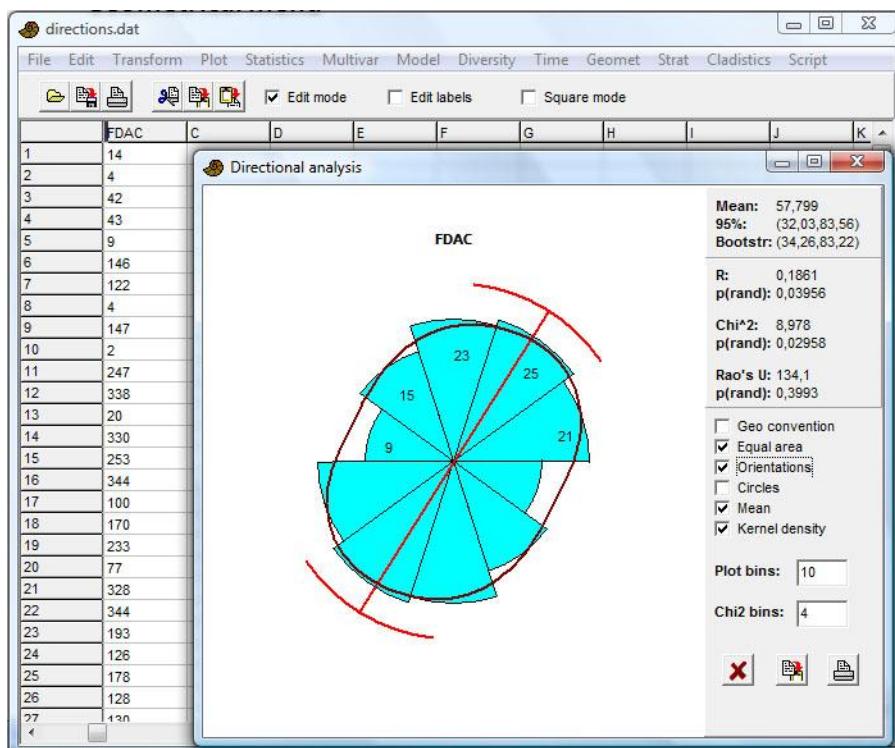
Laskar, J., P. Robutel, F. Joutel, M. Gastineau, A.C.M. Correia & B. Levrard. 2004. A long-term numerical solution for the insolation quantities of the Earth. *Astronomy & Astrophysics* 428:261-285.

Geometrical menu

Directions (one sample)

The module plots a rose diagram (polar histogram) of directions. Used for plotting current-oriented specimens, orientations of trackways, orientations of morphological features (e.g. terrace lines), etc.

One column of directional (0-360) or orientational (0-180) data in degrees is expected. Directional or periodic data in other forms (radians, 0-24 hours, etc.) must be converted to degrees using e.g. the Evaluate Expression module (Transform menu).



By default, the 'mathematical' angle convention of anticlockwise from east is chosen. If you use the 'geographical' convention of clockwise from north, tick the box.

You can also choose whether to have the abundances proportional to radius in the rose diagram, or proportional to area (equal area).

The "Kernel density" option plots a circular kernel density estimate.

The mean angle takes circularity into account. The 95 percent confidence interval on the mean is estimated according to Fisher (1983). It assumes circular normal distribution, and is not accurate for very large variances (confidence interval larger than 45 degrees) or small sample sizes. The bootstrapped 95% confidence interval on the mean uses 5000 bootstrap replicates. The graphic uses the bootstrapped confidence interval.

The R value (Rayleigh's spread) is given by:

$$R = \sqrt{\left(\sum_{i=1}^n \cos \theta_i \right)^2 + \left(\sum_{i=1}^n \sin \theta_i \right)^2}$$

R is further tested against a random distribution using Rayleigh's test for directional data (Davis 1986). Note that this procedure assumes evenly or unimodally distributed data - the test is not appropriate for bidirectional data. The p values are computed using an approximation given by Mardia (1972):

$$K = nR^2$$

$$p = e^{-K} \left(1 + \frac{2K - K^2}{4n} - \frac{24K - 132K^2 + 76K^3 - 9K^4}{288n^2} \right)$$

The Rao's spacing test (Batschelet 1981) for uniform distribution has test statistic

$$U = \frac{1}{2} \sum_{i=1}^n |T_i - \lambda|,$$

where $\lambda = 360^\circ/n$. $T_i = \theta_{i+1} - \theta_i$ for $i < n$, $T_n = 360^\circ - \theta_n + \theta_1$.

The p value is estimated by linear interpolation from the probability tables published by Russell & Levitin (1995).

A Chi-square test for uniform distribution is also available, with a user-defined number of bins (default 4).

The 'Orientations' option allows analysis of linear orientations (0-180 degrees). The Rayleigh test is then carried out by a directional test on doubled angles (this trick is described by Davis 1986); the Chi-square uses four bins from 0-180 degrees; the rose diagram mirrors the histogram around the origin.

References

Batschelet, E. 1981. Circular statistics in biology. Academic Press.

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Fisher, N.I. 1983. Comment on "A Method for Estimating the Standard Deviation of Wind Directions". *Journal of Applied Meteorology* 22:1971.

Mardia, K.V. 1972. Statistics of directional data. Academic Press.

Russell, G. S. & D.J. Levitin 1995. An expanded table of probability values for Rao's spacing test. *Communications in Statistics: Simulation and Computation* 24:879-888.

Directions (two samples)

Expects two columns of directional (0-360) or orientational (0-180) data in degrees.

Watson-Williams test for equal mean angle in two samples. With corrections due to Mardia (1972). The concentration parameter κ is maximum-likelihood, computed analytically. It should be larger than 1.0 for accurate testing. In addition, the test assumes similar angular variances (R values).

Reference

Mardia, K.V. 1972. Statistics of directional data. Academic Press.

Circular correlation

Testing for correlation between two directional or orientational variates. “Large N”. Two columns of directional (0-360) or orientational (0-180) data in degrees.

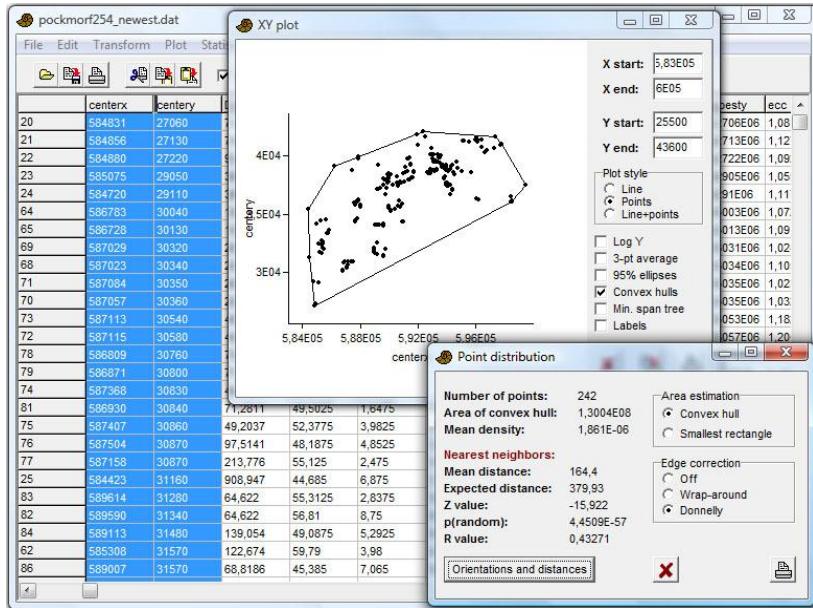
This module uses the circular correlation procedure and parametric significance test of Jammalamadaka & Sengupta (2001).

Reference

Jammalamadaka, S.R. & A. Sengupta. 2001. Topics in circular statistics. World Scientific.

Nearest neighbour point pattern analysis

This module tests for clustering or overdispersion of points given as two-dimensional coordinate values. The procedure assumes that elements are small compared to their distances, that the domain is predominantly convex, and $n>50$. Two columns of x/y positions are required. Applications of this module include spatial ecology (are in-situ brachiopods clustered), morphology (are trilobite tubercles overdispersed), and geology (distribution of e.g. volcanoes, earthquakes, springs).



The calculation of point distribution statistics using nearest neighbour analysis follows Davis 1986 with modifications. The area is estimated either by the smallest enclosing rectangle or using the convex hull, which is the smallest convex polygon enclosing the points. Both are inappropriate for points in very concave domains. Two different edge effect adjustment methods are available: wrap-around ("torus") and Donnelly's correction. Wrap-around edge detection is only appropriate for rectangular domains.

The probability that the distribution is random (Poisson process, giving an exponential nearest neighbour distribution) is presented, together with the R value:

$$R = \frac{2\bar{d}}{\sqrt{A/n}}$$

where \bar{d} is the observed mean distance between nearest neighbours, A is the area, and n is the number of points. Clustered points give $R<1$, Poisson patterns give $R\sim 1$, while overdispersed points give $R>1$.

The orientations (0-180 degrees) and lengths of lines between nearest neighbours, are also included. The orientations can be subjected to directional analysis to test whether the points are organised along lineaments (see Hammer 2009 for more advanced methods).

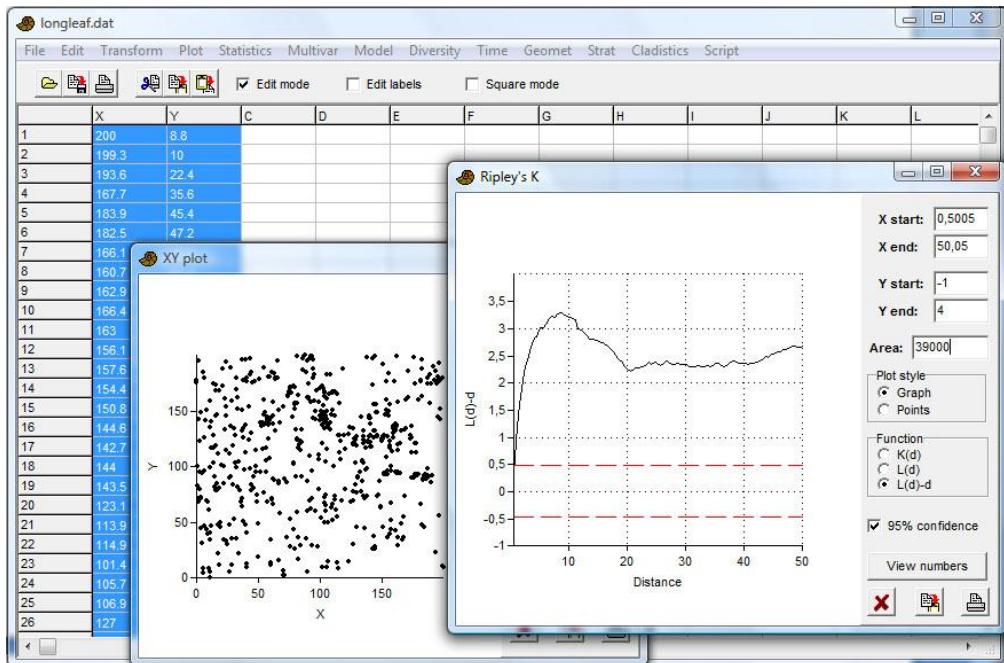
References

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Hammer, Ø. 2009. New methods for the statistical analysis of point alignments. Computers & Geosciences.

Ripley's K point pattern analysis

Ripley's K (Ripley 1979) is the average point density as a function of distance from every point. It is useful when point pattern characteristics change with scale, e.g. overdispersion over small distances but clustering over large distances. Two columns of x/y coordinates in a rectangular domain are expected.



Define the estimated intensity of the point pattern, with n points in an area A , as $\lambda = n/A$. The distance between points i and j is d_{ij} . The estimate of Ripley's K , as a function of distance, is then computed as

$$K(d) = \frac{1}{\lambda n} \sum_{i=1}^n \sum_{j \neq i} I(d_{ij} \leq d),$$

where the indicator function I is one if the argument is true, zero otherwise.

The normalization of K is such that for complete spatial randomness (CSR), $K(d)$ is expected to increase as the area of circles, i.e. $K(d) = \pi d^2$. The $L(d)$ function is a corresponding transformation of $K(d)$:

$$L(d) = \sqrt{\frac{K(d)}{\pi}}$$

For CSR, $L(d) = d$, and $L(d) - d = 0$. An approximate 95% confidence interval for $L(d) - d$ under CSR is given by $1.42\sqrt{A/n}$. Ripley's edge correction is included, giving weights to counts depending on the proportion of the test circle that is inside the rectangular domain.

The example above shows locations of trees in a forest. $L(d)-d$ is above the 95% confidence interval of CSR, indicating clustering. In addition, the spatial interactions seem most prominent at a scale of around 10 m, above which the curve flattens in the manner expected from CSR.

Area

For the correct calculation of Ripley's K , the area must be known. In the first run, the area is computed using the smallest bounding rectangle, but this can both over- and underestimate the real area. The area can therefore be adjusted by the user. An overestimated area will typically show up as a strong overall linear trend with positive slope for $L(d)-d$.

Fractal dimension

The fractal dimension (if any) can be estimated as the asymptotic linear slope in a log-log plot of $K(d)$. For CSR, the log-log slope should be 2.0. Fractals should have slopes less than 2.

References

Ripley, B.D. 1979. Tests of 'randomness' for spatial point patterns. *Journal of the Royal Statistical Society, ser. B* 41:368-374.

Gridding (spatial interpolation)

“Gridding” is the operation of spatial interpolation of scattered 2D data points onto a regular grid. Three columns with position (x,y) and corresponding data values are required.

Gridding allows the production of a map showing a continuous spatial estimate of some variate such as fossil abundance or thickness of a rock unit, based on scattered data points. The user can specify the size of the grid (number of rows and columns), but in the present version the spatial coverage of the map is generated automatically to cover the data points (the map will always be square).

A least-squares linear surface (trend) is automatically fitted to the data, removed prior to gridding and finally added back in. This is primarily useful for the semivariogram modelling and the kriging method.

Three algorithms are available:

Moving average

The value at a grid node is simply the average of the N closest data points, as specified by the user (the default is to use all data points). The points weighted in inverse proportion to distance. This algorithm is simple and will not always give good (smooth) results. One advantage is that the interpolated values will never exceed the range of the data points.

Thin-plate spline

Maximally smooth interpolator. Can overshoot in the presence of sharp bends in the surface.

Kriging

The user is required to specify a model for the semivariogram, by choosing one of three models (spherical, exponential or Gaussian) and corresponding parameters to fit the empirical semivariances as well as possible. The semivariogram is computed within each of a number of bins. Using the histogram option, choose a number of bins so that each bin (except possibly the rightmost ones) contains at least 30 distances. See e.g. Davis (1986) for more information.

The kriging procedure also provides an estimate of standard errors across the map (this depends on the semivariogram model being accurate). Kriging in PAST does not provide for anisotropic semivariance.

Reference

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Multivariate allometry

Finding and testing for allometry in a multivariate morphometric data set. A multivariate data set with variables (distance measurements) in columns, specimens in rows.

This method for investigating allometry in a multivariate data set is based on Jolicoeur (1963) with extensions by Kowalewski et al. (1997). The data are (automatically) log-transformed and subjected to PCA. The first principal component (PC1) is then regarded as a size axis (this is only valid if the variation accounted for by PC1 is large, say more than 80%). The allometric coefficient for each original variable is estimated by dividing the PC1 loading for that variable by the mean PC1 loading over all variables.

95% confidence intervals for the allometric coefficients are estimated by bootstrapping specimens. 2000 bootstrap replicates are made.

Missing data is supported by column average substitution.

References

Jolicoeur, P. 1963. The multivariate generalization of the allometry equation. *Biometrics* 19:497-499.

Kowalewski, M., E. Dyreson, J.D. Marcot, J.A. Vargas, K.W. Flessa & D.P. Hallmann. 1997. Phenetic discrimination of biometric simpletons: paleobiological implications of morphospecies in the lingulide brachiopod *Glottidia*. *Paleobiology* 23:444-469.

Fourier shape (2D)

Analysis of fossil outline shape (2D). Shape expressible in polar coordinates, sufficient number of digitized points to capture features. Digitized x/y coordinates around an outline. Specimens in rows, coordinates of alternating x and y values in columns (see Procrustes fitting in the Transform menu).

Accepts X-Y coordinates digitized around an outline. More than one shape (row) can be simultaneously analyzed. Points do not need to be totally evenly spaced. The shape must be expressible as a unique function in polar co-ordinates, that is, any straight line radiating from the centre of the shape must cross the outline only once.

The algorithm follows Davis (1986). The origin for the polar coordinate system is found by numerical approximation to the centroid. 128 points are then produced at equal angular increments around the outline, through linear interpolation. The centroid is then re-computed, and the radii normalized (size is thus removed from the analysis). The cosine and sine components are given for the first twenty harmonics, but note that only $N/2$ harmonics are 'valid', where N is the number of digitized points. The coefficients can be copied to the main spreadsheet for further analysis (e.g. by PCA).

The 'Shape view' window allows graphical viewing of the Fourier shape approximation(s).

Reference

Davis, J.C. 1986. Statistics and Data Analysis in Geology. John Wiley & Sons.

Elliptic Fourier shape analysis

Digitized x/y coordinates around an outline. Specimens in rows, coordinates of alternating x and y values in columns (see Procrustes fitting in the Transform menu).

More than one shape (row) can be simultaneously analyzed.

Elliptic Fourier shape analysis is in several respects superior to simple Fourier shape analysis. One advantage is that the algorithm can handle complicated shapes which may not be expressible as a unique function in polar co-ordinates. Elliptic Fourier shapes is now a standard method of outline analysis. The algorithm used in PAST is described by Ferson et al. (1985).

Cosine and sine components of x and y increments along the outline for the first 30 harmonics are given, but only the first $N/2$ harmonics should be used, where N is the number of digitized points. Size and positional translation are normalized away, and do not enter in the coefficients. However, no attempt is made to standardize rotation or starting point, so all specimens should be measured in a standard orientation. The coefficients can be copied to the main spreadsheet for further analysis (e.g. by PCA).

The 'Shape view' window allows graphical viewing of the elliptic Fourier shape approximation(s).

Reference

Ferson, S.F., F.J. Rohlf & R.K. Koehn. 1985. Measuring shape variation of two-dimensional outlines. *Systematic Zoology* 34:59-68.

Eigenshape analysis

Digitized x/y coordinates around an outline. Specimens in rows, coordinates of alternating x and y values in columns (see Procrustes fitting in the Transform menu).

Eigenshapes are principal components of outlines. The scatter plot of outlines in principal component space can be shown, and linear combinations of the eigenshapes themselves can be visualized.

The implementation in PAST is partly based on MacLeod (1999). It finds the optimal number of equally spaced points around the outline using an iterative search, so the original points need not be equally spaced. The eigenanalysis is based on the covariance matrix of the non-normalized turning angle increments around the outlines. The algorithm does not assume a closed curve, and the endpoints are therefore not constrained to coincide in the reconstructed shapes. Landmark-registered eigenshape analysis is not included. All outlines must start at the 'same' point.

Reference

MacLeod, N. 1999. Generalizing and extending the eigenshape method of shape space visualization and analysis. *Paleobiology* 25:107-138.

Thin-plate splines and warps

Digitized x/y landmark coordinates. Specimens in rows, coordinates of alternating x and y values in columns. Procrustes standardization recommended.

The first specimen (first row) is taken as a reference, with an associated square grid. The warps from this to all other specimens can be viewed. You can also choose the mean shape as the reference.

The 'Expansion factors' option will display the area expansion (or contraction) factor around each landmark in yellow numbers, indicating the degree of local growth. This is computed using the Jacobian of the warp. Also, the expansions are colour-coded for all grid elements, with green for expansion and purple for contraction.

At each landmark, the principal strains can also be shown, with the major strain in black and minor strain in brown. These vectors indicate directional stretching.

A description of thin-plate spline transformation grids is given by Dryden & Mardia (1998).

Partial warps

From the thin-plate spline window, you can choose to see the partial warps for a particular spline deformation. The first partial warp will represent some long-range (large scale) deformation of the grid, while higher-order warps will normally be connected with more local deformations. The affine component of the warp (also known as zeroth warp) represents linear translation, scaling, rotation and shearing. In the present version of PAST you can not view the principal warps.

When you increase the amplitude factor from zero, the original landmark configuration and a grid will be progressively deformed according to the selected partial warp.

Partial warp scores

From the thin-plate spline window, you can also choose to see the partial warp scores of all the specimens. Each partial warp score has two components (x and y), and the scores are therefore presented in scatter plots.

Reference

Dryden, I.L. & K.V. Mardia 1998. Statistical Shape Analysis. Wiley.

Relative warps

Ordination of a set of shapes. Digitized x/y landmark coordinates. Specimens in rows, coordinates of alternating x and y values in columns. Procrustes standardization recommended.

The relative warps can be viewed as the principal components of the set of thin-plate transformations from the mean shape to each of the shapes under study. It provides an alternative to direct PCA of the landmarks (see Shape PCA above).

The parameter *alpha* can be set to one of three values:

- *alpha=-1* emphasizes small-scale variation.
- *alpha=0* is PCA of the landmarks directly, and is equivalent to Shape PCA (see above) *but without including the affine (uniform) component*.
- *alpha=1* emphasizes large-scale variation.

The relative warps are ordered according to importance, and the first and second warps are usually the most informative. Note that the percentage values of the eigenvalues are relative to the total non-affine part of the transformation - the affine part is not included (see Shape PCA for relative warps with the affine component included).

The relative warps are visualized with thin-plate spline transformation grids. When you increase or decrease the amplitude factor away from zero, the original landmark configuration and grid will be progressively deformed according to the selected relative warp.

The relative warp scores of pairs of consecutive relative warps can be shown in scatter plots, and all scores can be shown in a numerical matrix.

The algorithm for computing the relative warps is taken from Dryden & Mardia (1998).

Reference

Dryden, I.L. & K.V. Mardia 1998. Statistical Shape Analysis. Wiley.

Size from landmarks (2D or 3D)

Digitized x/y or x/y/z landmark coordinates. Specimens in rows, coordinates with alternating x and y (and z for 3D) values in columns. Must not be Procrustes fitted or normalized for size!

Calculates the centroid size for each specimen (Euclidean norm of the distances from all landmarks to the centroid).

The values in the 'Normalized' column are centroid sizes divided by the square root of the number of landmarks - this might be useful for comparing specimens with different numbers of landmarks.

Normalize size

The 'Normalize size' option in the Transform menu allows you to remove size by dividing all coordinate values by the centroid size for each specimen. For 2D data you may instead use Procrustes coordinates, which are also normalized with respect to size.

See Dryden & Mardia (1998), p. 23-26.

Reference

Dryden, I.L. & K.V. Mardia 1998. Statistical Shape Analysis. Wiley.

Distance from landmarks (2D or 3D)

Digitized x/y or x/y/z landmark coordinates. Specimens in rows, coordinates with alternating x and y (and z for 3D) values in columns. May or may not be Procrustes fitted or normalized for size.

Calculates the Euclidean distances between two fixed landmarks for one or many specimens. You must choose two landmarks - these are named according to the name of the first column for the landmark (x value).

All distances from landmarks (EDMA)

Digitized x/y or x/y/z landmark coordinates. Specimens in rows, coordinates with alternating x and y (and z for 3D) values in columns. May or may not be Procrustes fitted or normalized for size.

This function will replace the landmark data in the data matrix with a data set consisting of distances between all pairs of landmarks, with one specimen per row. The number of pairs is $N(N-1)/2$ for N landmarks. This transformation will allow multivariate analysis of distance data, which are not sensitive to rotation or translation of the original specimens, so a Procrustes fitting is not mandatory before such analysis. Using distance data also allows log-transformation, and analysis of fit to the allometric equation for pairs of distances.

Missing data is supported by column average substitution.

Landmark linking

This function in the Geomet menu allows the selection of any pairs of landmarks to be linked with lines in the morphometric plots (thin-plate splines, partial and relative warps, etc.), to improve readability. The landmarks must be present in the main spreadsheet before links can be defined.

Pairs of landmarks are selected or deselected by clicking in the symmetric matrix. The set of links can also be saved in a text file. Note that there is little error checking in this module.

Strat menu

Unitary Associations

Unitary Associations analysis (Guex 1991) is a method for biostratigraphical correlation (see Angiolini & Bucher 1999 for a typical application). The data input consists of a presence/absence matrix with samples in rows and taxa in columns. Samples belonging to the same section (locality) must be assigned the same color, and ordered stratigraphically within each section such that the lowermost sample enters in the lowest row. Colors can be re-used in data sets with large numbers of sections .

Overview of the method

The method of Unitary Associations is logical, but rather complicated, consisting of a number of steps. For details, see Guex (1991). The implementation in PAST includes most of the features found in the original program, called BioGraph (Savary & Guex 1999), and thanks to a fruitful co-operation with Jean Guex it also includes a number of options and improvements not found in the present version of that program.

The basic idea is to generate a number of assemblage zones (similar to 'Oppel zones') which are optimal in the sense that they give maximal stratigraphic resolution with a minimum of superpositional contradictions. One example of such a contradiction would be a section containing a species A above a species B, while assemblage 1 (containing species A) is placed below assemblage 2 (containing species B). PAST (and BioGraph) carries out the following steps:

1. Residual maximal horizons

The method makes the range-through assumption, meaning that taxa are considered to have been present at all levels between the first and last appearance in any section. Then, any samples with a set of taxa that is contained in another sample are discarded. The remaining samples are called *residual maximal horizons*. The idea behind this throwing away of data is that the absent taxa in the discarded samples may simply not have been found even though they originally existed. Absences are therefore not as informative as presences.

2. Superposition and co-occurrence of taxa

Next, all pairs (A,B) of taxa are inspected for their superpositional relationships: A below B, B below A, A together with B, or unknown. If A occurs below B in one locality and B below A in another, they are considered to be co-occurring although they have never actually been found together.

The superpositions and co-occurrences of taxa can be viewed in the *biostratigraphic graph*. In this graph, taxa are coded as numbers. Co-occurrences between pairs of taxa are shown as solid blue lines. Superpositions are shown as dashed red lines, with long dashes from the above-occurring taxon and short dashes from the below-occurring taxon.

Some taxa may occur in so-called *forbidden sub-graphs*, which indicate inconsistencies in their superpositional relationships. Two of the several types of such sub-graphs can be plotted in PAST: C_n *cycles*, which are superpositional cycles (A->B->C->A), and S_3 *circuits*, which are inconsistencies of the type 'A co-occurring with B, C above A, and C below B'. Interpretations of such forbidden sub-graphs are suggested by Guex (1991).

3. Maximal cliques

Maximal cliques are groups of co-occurring taxa not contained in any larger group of co-occurring taxa. The maximal cliques are candidates for the status of unitary associations, but will be further processed below. In PAST, maximal cliques receive a number and are also named after a maximal horizon in the original data set which is identical to, or contained in (marked with asterisk), the maximal clique.

4. Superposition of maximal cliques

The superpositional relationships between maximal cliques are decided by inspecting the superpositional relationships between their constituent taxa, as computed in step 2. Contradictions (some taxa in clique A occur below some taxa in clique B, and vice versa) are resolved by a 'majority vote'. The contradictions between cliques can be viewed in PAST.

The superpositions and co-occurrences of cliques can be viewed in the *maximal clique graph*. In this graph, cliques are coded as numbers. Co-occurrences between pairs of cliques are shown as solid blue lines. Superpositions are shown as dashed red lines, with long dashes from the above-occurring clique and short dashes from the below-occurring clique. Also, cycles between maximal cliques (see below) can be viewed as green lines.

5. Resolving cycles

It will sometimes be the case that maximal cliques are now ordered in cycles: A is below B, which is below C, which is below A again. This is clearly contradictory. The 'weakest link' (superpositional relationship supported by fewest taxa) in such cycles is destroyed.

6. Reduction to unique path

At this stage, we should ideally have a single path (chain) of superpositional relationships between maximal cliques, from bottom to top. This is however often not the case, for example if A and B are below C, which is below D, or if we have isolated paths without any relationships (A below B and C below D). To produce a single path, it is necessary to merge cliques according to special rules.

7. Post-processing of maximal cliques

Finally, a number of minor manipulations are carried out to 'polish' the result: Generation of the 'consecutive ones' property, reinsertion of residual virtual co-occurrences and superpositions, and compaction to remove any generated non-maximal cliques. For details on these procedures, see Guex (1991). At last, we now have the Unitary Associations, which can be viewed in PAST.

The unitary associations have associated with them an index of similarity from one UA to the next, called D:

$$D_i = |UA_i - UA_{i-1}| / |UA_i| + |UA_{i-1} - UA_i| / |UA_{i-1}|$$

8. Correlation using the Unitary Associations

The original samples are now correlated using the unitary associations. A sample may contain taxa which uniquely places it in a unitary association, or it may lack key taxa which could differentiate between two or more unitary associations, in which case only a range can be given. These correlations can be viewed in PAST.

9. Reproducibility matrix

Some unitary associations may be identified in only one or a few sections, in which case one may consider to merge unitary associations to improve the geographical reproducibility (see below). The reproducibility matrix should be inspected to identify such unitary associations. A UA which is uniquely identified in a section is shown as a black square, while ranges of UAs (as given in the correlation list) are shown in gray.

10. Reproducibility graph and suggested UA merges (biozonation)

The reproducibility graph (G_k' in Guex 1991) shows the superpositions of unitary associations that are actually observed in the sections. PAST will internally reduce this graph to a unique maximal path (Guex 1991, section 5.6.3), and in the process of doing so it may merge some UAs. These mergers are shown as red lines in the reproducibility graph. The sequence of single and merged UAs can be viewed as a suggested biozonation.

Special functionality

The implementation of the Unitary Associations method in PAST includes a number of options and functions which have not yet been described in the literature. For questions about these, please contact us.

References

Angiolini, L. & H. Bucher. 1999. Taxonomy and quantitative biochronology of Guadalupian brachiopods from the Khuff Formation, Southeastern Oman. *Geobios* 32:665-699.

Guex, J. 1991. Biochronological Correlations. Springer Verlag.

Savary, J. & J. Guex. 1999. Discrete Biochronological Scales and Unitary Associations: Description of the BioGraph Computer Program. *Meemoires de Geologie (Lausanne)* 34.

Ranking-Scaling

Ranking-Scaling (Agterberg & Gradstein 1999) is a method for quantitative biostratigraphy based on *events* in a number of wells or sections. The data input consists of wells in rows with one well per row, and events (e.g. FADs and/or LADs) in columns. The values in the matrix are depths of each event in each well, increasing upwards (you may want to use negative values to achieve this). Absences are coded as zero. If only the order of events is known, this can be coded as increasing whole numbers (ranks, with possible ties for co-occurring events) within each well.

The implementation of ranking-scaling in PAST is not comprehensive, and advanced users are referred to the RASC and CASC programs of Agterberg and Gradstein.

Overview of the method

The method of Ranking-Scaling proceeds in two steps:

1. Ranking

The first step of Ranking-Scaling is to produce a single, comprehensive stratigraphic ordering of events, even if the data contains contradictions (event A over B in one well, but B over A in another), or longer cycles (A over B over C over A). This is done by 'majority vote', counting the number of times each event occurs above, below or together with all others. Technically, this is achieved by Presorting followed by the Modified Hay Method (Agterberg & Gradstein 1999).

2. Scaling

The biostratigraphic analysis may end with ranking, but additional insight may be gained by estimating stratigraphic distances between the consecutive events. This is done by counting the number of observed superpositional relationships (A above or below B) between each pair (A,B) of consecutive events. A low number of contradictions implies long distance.

Some computed distances may turn out to be negative, indicating that the ordering given by the ranking step was not optimal. If this happens, the events are re-ordered and the distances re-computed in order to ensure only positive inter-event distances.

RASC in PAST

Parameters

- Well threshold: The minimum number of wells in which an event must occur in order to be included in the analysis
- Pair threshold: The minimum number of times a relationship between events A and B must be observed in order for the pair (A,B) to be included in the ranking step

- Scaling threshold: Pair threshold for the scaling step
- Tolerance: Used in the ranking step (see Agterberg & Gradstein)
-

Ranking

The ordering of events after the ranking step is given, with the first event at the bottom of the list. The "Range" column indicates uncertainty in the position.

Scaling

The ordering of the events after the scaling step is given, with the first event at the bottom of the list. For an explanation of all the columns, see Agterberg & Gradstein (1999).

Event distribution

A plot showing the number of events in each well, with the wells ordered according to number of events.

Scattergrams

For each well, the depth of each event in the well is plotted against the optimum sequence (after scaling). Ideally, the events should plot in an ascending sequence.

Dendrogram

Plot of the distances between events in the scaled sequence, including a dendrogram which may aid in zonation.

Reference

Agterberg, F.P. & F.M. Gradstein. 1999. The RASC method for Ranking and Scaling of Biostratigraphic Events. In: Proceedings Conference 75th Birthday C.W. Drooger, Utrecht, November 1997. *Earth Science Review* 46(1-4):1-25.

CONOP

Table of depths/levels, with wells/sections in rows and event pairs in columns: FADs in odd columns and LADs in even columns. Missing events coded with zeros.

PAST includes a simple version of Constrained Optimization (Kemple et al. 1989). Both FAD and LAD of each taxon must be specified in alternate columns. Using so-called Simulated Annealing, the program searches for a global (composite) sequence of events that implies a minimal total amount of range extension (penalty) in the individual wells/sections. The parameters for the optimization procedure include an initial annealing temperature, the number of cooling steps, the cooling ratio (percentage lower than 100), and the number of trials per step. For explanation and recommendations, see Kemple et al. (1989).

Output windows include the optimization history with the temperature and penalty as function of cooling step, the global composite solution and the implied ranges in each individual section.

The implementation of CONOP in PAST is based on a FORTRAN optimization core provided by Sadler and Kemple.

Reference

Kemple, W.G., P.M. Sadler & D.J. Strauss. 1989. A prototype constrained optimization solution to the time correlation problem. In Agterberg, F.P. & G.F. Bonham-Carter (eds), Statistical Applications in the Earth Sciences. Geological Survey of Canada Paper 89-9:417-425.

Appearance Event Ordination

Appearance Event Ordination (Alroy 1994, 2000) is a method for biostratigraphical seriation and correlation. The data input is in the same format as for Unitary Associations, consisting of a presence/absence matrix with samples in rows and taxa in columns. Samples belonging to the same section (locality) must be assigned the same color, and ordered stratigraphically within each section such that the lowermost sample enters in the lowest row. Colors can be re-used in data sets with large numbers of sections.

The implementation in PAST is based on code provided by John Alroy. It includes Maximum Likelihood AEO (Alroy 2000).

References

Alroy, J. 1994. Appearance event ordination: a new biochronologic method. *Paleobiology* 20:191-207.

Alroy, J. 2000. New methods for quantifying macroevolutionary patterns and processes. *Paleobiology* 26:707-733.

Diversity curve

Abundance or presence/absence matrix with samples in rows (lowest sample at bottom) and taxa in columns

Found in the 'Strat' menu, this simple tool allows plotting of diversity curves from occurrence data in a stratigraphical column. Note that samples should be in stratigraphical order, with the uppermost (youngest) sample in the uppermost row. Data are subjected to the range-through assumption (absences between first and last appearance are treated as presences). Originations and extinctions are in absolute numbers, not percentages.

The 'Endpoint correction' option counts a FAD or LAD in a sample as 0.5 instead of 1 in that sample. Both FAD and LAD in the sample counts as 0.33.

Range confidence intervals

Estimation of confidence intervals for first or last appearances or total range, for one taxon.

Random distribution of fossiliferous horizons through the stratigraphic column or through time. Section should be continuously sampled.

The number of horizons containing the taxon, and levels or dates of first and last occurrences of the taxon.

Assuming a random (Poisson) distribution of fossiliferous horizons, confidence intervals for the stratigraphic range of one taxon can be calculated given the first occurrence datum (level), last occurrence datum, and total number of horizons where the taxon is found (Strauss & Sadler 1989, Marshall 1990).

No data are needed in the spreadsheet. The program will ask for the number of horizons where the taxon is found, and levels or dates for the first and last appearances. If necessary, use negative values to ensure that the last appearance datum has a higher numerical value than the first appearance datum. 80%, 95% and 99% confidence intervals are calculated for the FAD considered in isolation, the LAD considered in isolation, and the total range. The value *alpha* is the length of the confidence interval divided by the length of the observed range.

Be aware that the assumption of random distribution will not hold in many real situations.

References

Marshall, C.R. 1990. Confidence intervals on stratigraphic ranges. *Paleobiology* 16:1-10.

Strauss, D. & P.M. Sadler. 1989. Classical confidence intervals and Bayesian probability estimates for ends of local taxon ranges. *Mathematical Geology* 21:411-427.

Distribution-free range confidence intervals

Estimation of confidence intervals for first or last appearances. Assumes no correlation between stratigraphic position and gap size. Section should be continuously sampled. One column per taxon, with levels or dates of all horizons where the taxon is found.

This method (Marshall 1994) does not assume random distribution of fossiliferous horizons. It requires that the levels or dates of all horizons containing the taxon are given.

The program outputs upper and lower bounds on the lengths of the confidence intervals, using a 95 percent confidence probability, for confidence levels of 50, 80 and 95 percent. Values which can not be calculated are marked with an asterisk (see Marshall 1994).

Reference

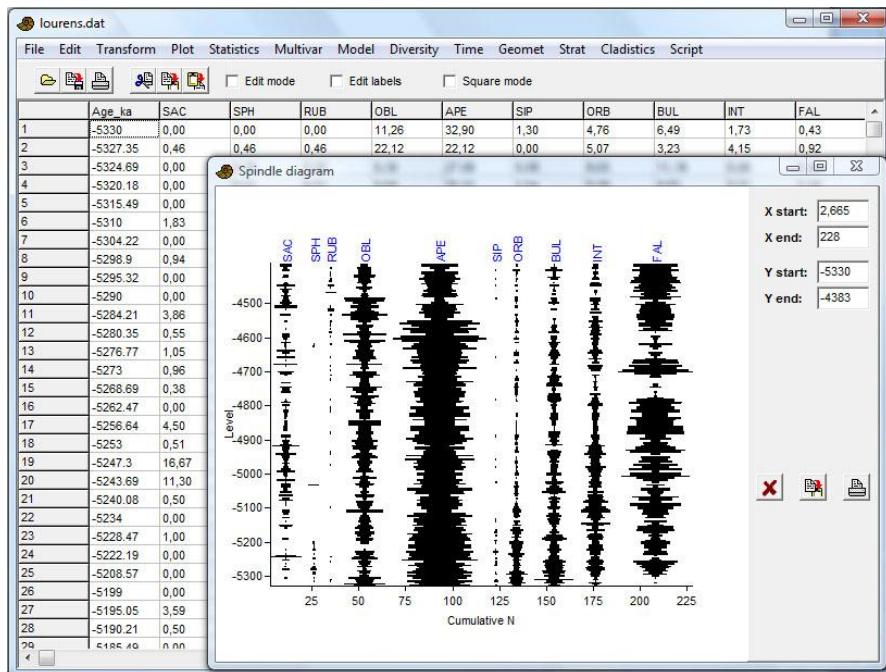
Marshall, C.R. 1994. Confidence intervals on stratigraphic ranges: partial relaxation of the assumption of randomly distributed fossil horizons. *Paleobiology* 20:459-469.

Spindle diagram

Standard plot type used in paleontology to illustrate abundance of fossil taxa through a stratigraphic section or core. Samples are entered in rows, taxa in columns. The program will ask whether the first column contains stratigraphic levels (e.g. meters or years).

If stratigraphic levels are given, each box will be drawn from the given level to the level in the next row. Therefore a last, extra “dummy” row should be given, with a final stratigraphic level but zero for all the taxa. If levels are given in depths or ages, negative numbers should be used in order to ensure correct orientation of the figure.

If stratigraphic levels are not given, all boxes will be given equal height. The topmost sample should be entered in the first row.



Filter events

text

Cladistics

Parsimony analysis

Warning: the cladistics package in PAST is fully operational, but lacking in comprehensive functionality. The heuristic algorithms seem not to perform quite as well as in some other programs (this is being looked into). The PAST cladistics package is adequate for education and initial data exploration, but for more 'serious' work we recommend a specialized program such as PAUP.

Semi-objective analysis of relationships between taxa from morphological or genetic evidence.

Character matrix with taxa in rows, outgroup in first row. For calculation of stratigraphic congruence indices, first and last appearance datums must be given in the first two columns.

Algorithms are from Kitching et al. (1998).

Character states should be coded using integers in the range 0 to 255, or with the letters c, a, g, t, u (upper or lower case). The first taxon is treated as the outgroup, and will be placed at the root of the tree.

Missing values are coded with a question mark (?) or the value -1. Please note that PAST does not collapse zero-length branches. Because of this, missing values can lead to a proliferation of equally shortest trees *ad nauseam*, many of which are in fact equivalent.

There are four algorithms available for finding short trees:

Branch-and-bound

The branch-and-bound algorithm is guaranteed to find all shortest trees. The total number of shortest trees is reported, but a maximum of 10000 trees are saved. The branch-and-bound algorithm can be very time consuming for data sets with more than 12 taxa.

Exhaustive

The exhaustive algorithm evaluates all possible trees. Like the branch-and-bound algorithm it will necessarily find all shortest trees, but it is very slow. For 12 taxa, more than 600 million trees are evaluated! The only advantage over branch-and-bound is the plotting of tree length distribution. This histogram may indicate the 'quality' of your matrix, in the sense that there should be a tail to the left such that few short trees are 'isolated' from the greater mass of longer trees (but see Kitching et al. 1998 for critical comments on this). For more than 8 taxa, the histogram is based on a subset of tree lengths and may not be accurate.

Heuristic, nearest neighbour interchange

This heuristic algorithm adds taxa sequentially in the order they are given in the matrix, to the branch where they will give least increase in tree length. After each taxon is added, all nearest neighbour trees are swapped to try to find an even shorter tree.

Like all heuristic searches, this one is much faster than the algorithms above and can be used for large numbers of taxa, but is not guaranteed to find all or any of the most parsimonious trees. To decrease the likelihood of ending up on a suboptimal local minimum, a number of *reorderings* can be specified. For each reordering, the order of input taxa will be randomly permuted and another heuristic search attempted.

Please note: Because of the random reordering, the trees found by the heuristic searches will normally be different each time. To reproduce a search exactly, you will have to start the parsimony module again from the menu, using the same value for "Random seed". This will reset the random number generator to the seed value.

Heuristic, subtree pruning and regrafting

This algorithm (SPR) is similar to the one above (NNI), but with a more elaborate branch swapping scheme: A subtree is cut off the tree, and regrafting onto all other branches in the tree is attempted in order to find a shorter tree. This is done after each taxon has been added, and for all possible subtrees. While slower than NNI, SPR will often find shorter trees.

Heuristic, tree bisection and reconnection

This algorithm (TBR) is similar to the one above (SPR), but with an even more complete branch swapping scheme. The tree is divided into two parts, and these are reconnected through every possible pair of branches in order to find a shorter tree. This is done after each taxon is added, and for all possible divisions of the tree. TBR will often find shorter trees than SPR and NNI, at the cost of longer computation time.

Character optimization criteria

Three different optimization criteria are available:

Wagner

Characters are reversible and ordered, meaning that 0->2 costs more than 0->1, but has the same cost as 2->0.

Fitch

Characters are reversible and unordered, meaning that all changes have equal cost. This is the criterion with fewest assumptions, and is therefore generally preferable.

Dollo

Characters are ordered, but acquisition of a character state (from lower to higher value) can happen only once. All homoplasy is accounted for by secondary reversals. Hence, 0->1 can only happen once, normally relatively close to the root of the tree, but 1->0 can happen any number of times further up in the tree. (This definition has been debated on the PAST mailing list, especially whether Dollo characters need to be ordered).

Bootstrap

Bootstrapping is performed when the 'Bootstrap replicates' value is set to non-zero. The specified number of replicates (typically 100 or even 1000) of your character matrix are made, each with randomly weighted characters. The bootstrap value for a group is the percentage of replicates supporting that group. A replicate supports the group if the group exists in the majority rule consensus tree of the shortest trees made from the replicate.

Warning: Specifying 1000 bootstrap replicates will clearly give a thousand times longer computation time than no bootstrap! Exhaustive search with bootstrapping is unrealistic and is not allowed.

Cladogram plotting

All shortest (most parsimonious) trees can be viewed, up to a maximum of 10000 trees. If bootstrapping has been performed, a bootstrap value is given at the root of the subtree specifying each group.

Character states can also be plotted onto the tree, as selected by the 'Character' buttons. This character reconstruction is unique only in the absence of homoplasy. In case of homoplasy, character changes are placed as close to the root as possible, favouring one-time acquisition and later reversal of a character state over several independent gains (so-called *accelerated transformation*).

The 'Phylogram' option allows plotting of trees where the length of vertical lines (joining clades) is proportional to branch length.

Consistency index

The per-character consistency index (ci) is defined as m/s , where m is the minimum possible number of character changes (steps) on any tree, and s is the actual number of steps on the current tree. This index hence varies from one (no homoplasy) and down towards zero (a lot of homoplasy). The ensemble consistency index CI is a similar index summed over all characters.

Retention index

The per-character retention index (ri) is defined as $(g-s)/(g-m)$, where m and s are as for the consistency index, while g is the maximal number of steps for the character on any cladogram (Farris 1989). The retention index measures the amount of synapomorphy on the tree, and varies from 0 to 1.

Please note that in the present version, the retention index is only correctly calculated when using Fitch optimisation.

Consensus tree

The consensus tree of all shortest (most parsimonious) trees can also be viewed. Two consensus rules are implemented: Strict (groups must be supported by all trees) and majority (groups must be supported by more than 50% of the trees).

Bremer support (decay index)

The Bremer support for a clade is the number of extra steps you need to construct a tree (consistent with the characters) where that clade is no longer present. There are reasons to prefer this index rather than the bootstrap value. PAST does not compute Bremer supports directly, but for smaller data sets it can be done 'manually' as follows:

- Perform parsimony analysis with exhaustive search or branch-and-bound. Take note of the clades and the length N of the shortest tree(s) (for example 42). If there are more than one shortest tree, look at the strict consensus tree. Clades which are no longer found in the consensus tree have a Bremer support value of 0.
- In the box for 'Longest tree kept', enter the number $N+1$ (43 in our example) and perform a new search.
- Additional clades which are no longer found in the strict consensus tree have a Bremer support value of 1.
- For 'Longest tree kept', enter the number $N+2$ (44) and perform a new search. Clades which now disappear in the consensus tree have a Bremer support value of 2.
- Continue until all clades have disappeared.

Stratigraphic congruence indices

For calculation of stratigraphic congruence indices, the first two columns in the data matrix must contain the first and last appearance datums, respectively, for each taxon. These datums must be given such that the youngest age (or highest stratigraphic level) has the highest numerical value. You may need to use negative values to achieve this (e.g. 400 million years before present is coded as -400.0). The box "FADs/LADs in first columns" in the Parsimony dialogue must be ticked.

The Stratigraphic Congruence Index (SCI) of Huelsenbeck (1994) is defined as the proportion of stratigraphically consistent nodes on the cladogram, and varies from 0 to 1. A node is stratigraphically consistent when the oldest first occurrence above the node is the same age or younger than the first occurrence of its sister taxon (node).

The Relative Completeness Index (RCI) of Benton & Storrs (1994) is defined as $(1 - \text{MIG}/\text{SRL}) \times 100\%$, where MIG (Minimum Implied Gap) is the sum of the durations of ghost ranges and SRL is the sum of the durations of observed ranges. The RCI can become negative, but will normally vary from 0 to 100.

The Gap Excess Ratio (GER) of Wills (1999) is defined as $1 - (\text{MIG} - \text{G}_{\min}) / (\text{G}_{\max} - \text{G}_{\min})$ where G_{\min} is the minimum possible sum of ghost ranges on any tree (that is, the sum of distances between successive FADs), and G_{\max} is the maximum (that is, the sum of distances from first FAD to all other FADs).

These indices are further subjected to a permutation test, where all dates are randomly redistributed on the different taxa 1000 times. The proportion of permutations where the recalculated index exceeds the original index is given. If small (e.g. $p < 0.05$), this indicates a statistically significant departure from the null hypothesis of no congruency between cladogram and stratigraphy (in other words, you have significant congruency). The permutation probabilities of RCI and GER are equal for any given set of permutations, because they are based on the same value for MIG.

References

Benton, M.J. & G.W. Storrs. 1994. Testing the quality of the fossil record: paleontological knowledge is improving. *Geology* 22:111-114.

Farris, J.S. 1989. The retention index and the rescaled consistency index. *Cladistics* 5:417-419.

Huelsenbeck, J.P. 1994. Comparing the stratigraphic record to estimates of phylogeny. *Paleobiology* 20:470-483.

Kitching, I.J., P.L. Forey, C.J. Humphries & D.M. Williams. 1998. *Cladistics*. Oxford University Press.

Wills, M.A. 1999. The gap excess ratio, randomization tests, and the goodness of fit of trees to stratigraphy. *Systematic Biology* 48:559-580.

Scripting

PAST includes a simple scripting (programming) language that allows you to make your own algorithms within the PAST framework. The language is a simplified Pascal, but with full matrix support and a library of mathematical, statistical and user-interface functions.

Important: As of the present version, this new scripting capability is rudimentary and not entirely stable. It will be rapidly improved in coming versions!

LANGUAGE STRUCTURE

Variables

Variable names can be of any length. There is no declaration of variables. All variables have global scope. The assignment operator is :=.

Comments

Upon encountering the comment character '#' in the source, the rest of the line is skipped.

begin ... end

Blocks are marked with begin ... end. The entire script must be contained within a begin end pair.

if ... else

Example:

```
if fad>0 begin
  lad:=fad+1;
  fad:=0;
end else lad:=0;
```

for ... to

Example:

```
for i:=1 to 10 a[i]:=0;
```

while

Example:

```
while a[i]=0 i:=i+1;
```

procedure

Procedures have no arguments, return values or local variables. All communication with the procedure must go through global variables. A procedure call must contain a dummy argument. The procedure definition must be inside the outermost begin ... end pair. Example:

```
begin
  procedure writehello
  begin
    message("Hello");
  end;
```

```
writehello(0);
end;
```

Types

Four types are available: Double-precision numbers, vectors, arrays and strings. Types are implicitly defined at first usage. "Sensible" type conversions are performed automatically.

Examples:

```
a:=14.3;
b:="Hi!";
c:=array(5,10); c[3,7]:=a;
d:=vector(10); d[7]:=17;
```

Operators

Operator	Supported types
+	double+double, vector+vector, array+array, array+double, vector+double
-	double-double, vector-vector, array-array, array-double, vector-double
*	double*double, array*array (array multiplication), array*double, vector*double
/	double/divide, array/divide, vector/divide
^	double^double (power)
%	double%double (integer modulo)
&	string&string, string&double, double&string (concatenation)
=	double=double, string=string
>	double>double, string>string
<	double<string<>
>=	double>=double, string>=string
<=	double<=double, string<=string
<>	double<>double, string<>string
and	double and double
or	double or double

MATHEMATICAL FUNCTIONS

Most of these functions support double, vector and array types.

Function	Comments
abs(x)	Absolute value
cos(x)	Cosine
sin(x)	Sine
tan(x)	Tangent

exp(x)	Exponential, e^x
log(x)	Natural logarithm
sqrt(x)	Square root
int(x)	Round down to integer
rnd(x)	Random, uniformly distributed number in the range [0..x)
fprob(f, df1, df2)	p value for F distribution
tprob(t, df)	p value for Student's t distribution
zprob(z)	p value for standardized normal distribution (z test)
chi2prob(chi2, df)	p value for chi^2 distribution
anova(col1,col2)	F statistic for one-way ANOVA, columns col1-col2 from spreadsheet

Array and vector operations

Function	Comments
nrows(x)	Number of rows in vector or array x
ncols(x)	Number of columns in array x
array(m,n)	Allocate an array of size m,n
vector(m)	Allocate a vector of length m
row(x,m)	Row m in array x, as a vector
column(x,n)	Column n in array x, as a vector
inv(x)	Inverse of double, or square array
mean(x)	Mean value of vector or array
eigen(x)	For a square NxN matrix x, returns a Nx(N+1) matrix with N sorted eigenvectors in the first N columns, and the eigenvalues in the last column.
cov(x)	Variance-covariance matrix of array x
sort2(x1, x2)	Sort the vectors x1 and x2, on x1. Returns an array with the sorted vectors in the two columns.

Distance matrices

Note that all these functions return symmetric distance matrices. For similarity indices such as Jaccard, the complement $1-x$ is returned. A value -1 in the input matrix is treated as missing value.

Function	Comments
eucliddist(x)	Symmetric Euclidean distance matrix of array x
chorddist(x)	Chord distance
cosdist(x)	Complement of cosine similarity
dicedist(x)	Complement of Dice similarity
jaccarddist(x)	Complement of Jaccard similarity
morisitadist(x)	Complement of Morisita similarity
horndist(x)	Complement of Horn similarity
manhattandist(x)	Manhattan distance

corrdist(x)	Complement of correlation similarity
rhodist(x)	Complement of non-parametric correlation
raupcrickdist(x)	Complement of Raup-Crick similarity
hammingdist(x1, x2)	Hamming distance
simpsondist(x1, x2)	Complement of Simpson similarity

USER INTERFACE FUNCTIONS AND PROCEDURES

Functions:

Function	Comments
spreadsheetarray(0 1)	Returns the selected array of numbers in the spreadsheet. Call with argument 1 for replacing missing values with column averages.
spreadsheetcolumn(n)	Returns column n from the spreadsheet, as a vector
spreadsheetsymbols(0)	Returns a vector of integer values for symbols (colors) in the selected array.
selectedrowlabel(m)	Returns a string containing the row label of row m in the selected array.
selectedrowlabel(n)	Returns a string containing the column label of column n in the selected array.

Procedures:

message(x)	Displays a message dialogue showing x (string or double)
setspreadsheet(m, n, x)	Set the contents of cell m,n in spreadsheet to x (string or double)
opennumwindow(m, n)	Opens a numerical output (table) window with m rows and n columns.
setnumwindow(m, n, x)	Set the contents of cell m,n in the numerical output window to x (string or double)

Graphics

Only one graphics window is available at any one time. Before displaying graphics, the window must be opened using the procedure openwindow. Coordinates can be in any unit, as the window is automatically scaled to accomodate its contents.

Colors are defined as 24-bit integers (R, G, B has 8 bits each), or using the pre-defined constants colred, colblack, colblue, colyellow, colgreen, colpurple.

Procedures:

openwindow	Opens the graphics window
drawline(x1, y1, x2, y2)	Draws a blue line.
drawstring(x1, y1, string, color)	Draws a text string.
drawvectorpoints(x, color)	Draws a set of points contained in the vector x, with indices on the horizontal axis.

drawxypoints(x, y, color)	Draws a set of points with x,y coordinates contained in the vectors x and y.
drawxysymbols(x, y, symbols)	Draws a set of points with x,y coordinates contained in the vectors x and y. The symbols vector contains integer values for symbols (colors).
drawhistogram(x, nbins, color)	Draws a histogram of the values contained in the vector x, with the number of bins given.

EXAMPLES

1. Take the first two columns of the spreadsheet, and display a log-log scatter diagram:

```

begin
  x:=spreadsheetcolumn(1);
  y:=spreadsheetcolumn(2);
  openwindow;
  drawxypoints(log(x), log(y), colblack);
end;

```

2. Carry out a Principal Components Analysis, with graphical and numerical output:

```

begin
  data:=spreadsheetarray(1);
  eig:=eigen(cov(data));
  eigvectors:=array(nrows(eig), ncols(eig)-1);
  for i:=1 to nrows(eig) for j:=1 to ncols(eig)-1
    eigvectors[i,j]:=eig[i,j];
  scores:=data*eigvectors;

  openwindow;
  drawxysymbols(column(scores,1), column(scores,2), spreadsheetsymbols(0));
  for i:=1 to nrows(scores)
    drawstring(scores[i,1], scores[i,2], spreadsheetrowlabel(i), colblack);

  opennumwindow(nrows(scores)+1, ncols(scores)+1);
  for i:=1 to nrows(scores) for j:=1 to ncols(scores)
    setnumwindow(i+1, j+1, scores[i,j]);

  for i:=1 to nrows(scores) setnumwindow(i+1, 1, spreadsheetrowlabel(i));
  for i:=1 to ncols(scores) setnumwindow(1, i+1, "Axis "&i);
end;

```