

pROC: display and analyze ROC curves

Tools for visualizing, smoothing and comparing receiver operating characteristic (ROC curves). (Partial) area under the curve (AUC) can be compared with statistical tests based on U-statistics or bootstrap. Confidence intervals can be computed for (p)AUC or ROC curves.

1. Installing and loading pROC

In the S+ command prompt, type:

```
install.pkgutils()  
library(pkgutils)
```

Download the “binary” file for your platform from <http://expasy.org/pROC/>, and save it in your S+ working directory. Type:

```
install.packages("filename.ext", repos=NULL)
```

taking care to replace filename.ext by the name of the downloaded file.

After successful installation, the package must be loaded through the File > Load Library...

More details about installation can be found on <http://expasy.org/pROC/-installS.html>.

2. Using pROC

pROC is available from the Statistics menu of the main menu bar of S+. The options are:

- ROC curve: to create and plot a ROC curve, and compute confidence intervals and (partial) AUC.
- ROC curve comparison: statistical tests for two ROC curves.
- Help: this help file

2.1 Abbreviations

The following abbreviations are employed extensively in this package:

- ROC: receiver operating characteristic
- AUC: area under the ROC curve
- pAUC: partial area under the ROC curve
- CI: confidence interval
- SP: specificity
- SE: sensitivity

3. ROC curve

The ROC curve dialog allows to build a ROC curve, and to compute its (partial) area under the curve, and confidence intervals.

3.1 General tab

The screenshot shows the 'Build a ROC curve' dialog box with the 'General' tab selected. The dialog has five tabs: General, Smoothing, AUC, CI, and Plot. The 'General' tab contains the following sections:

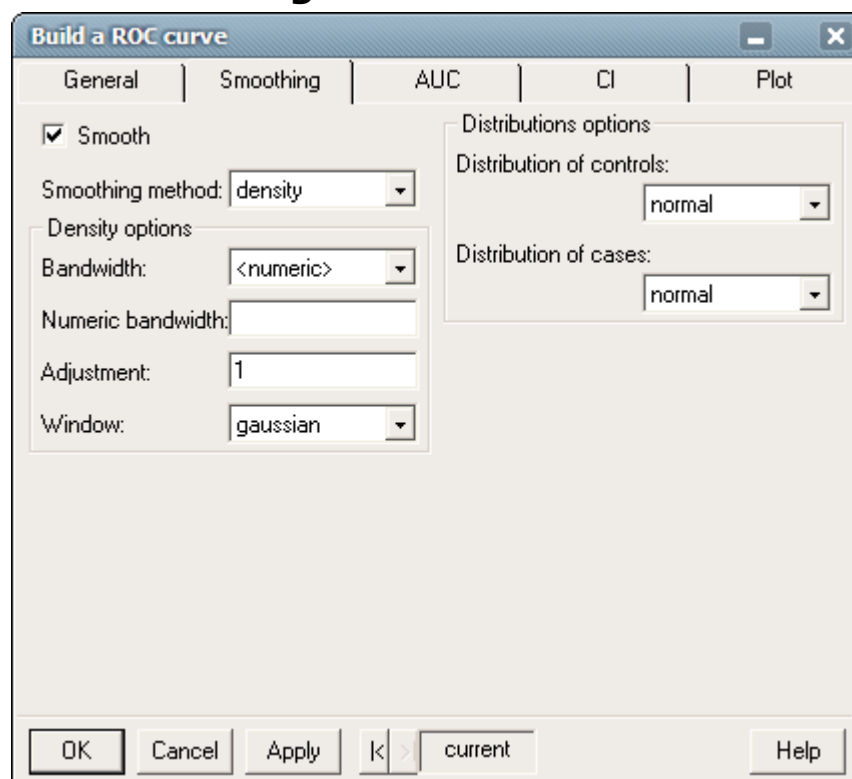
- Data Selection:**
 - Data Set: SubArachnoidH
 - ☒ Remove NAs
 - Test variable: s100b
- State:**
 - State variable: outcome
 - Level for controls: Good
 - Level for cases: Poor
- Actions:**
 - ☒ Smooth
 - ☒ Compute Area Under Curve (AUC)
 - ☒ Compute Confidence Interval (CI)
 - ☒ Plot
- Options:**
 - ☒ Percent
 - Direction: auto
- Output:**
 - ☒ Report
 - Coordinates: <custom>
 - Custom threshold(s):
- Save ROC curve:**
 - Save as: last.ROC

At the bottom, there are buttons for OK, Cancel, Apply, a navigation button with left and right arrows, a 'current' label, and a Help button.

- Data Selection
 - Data Set: the name of the data frame to use. The dataset must be loaded into S+ (for example with the File > Import data menu).
 - Remove NAs: if the Test or State variable contain missing values, remove the whole row of the data.
 - Test variable: the name of the test (continuous) variable to test within the data set. It must be a column of the selected data set.
- Actions
 - Smooth: check to apply smoothing to the ROC curve. More details can be set in the Smoothing tab.
 - Compute Area Under Curve (AUC): uncheck if you do not want to compute the AUC. More details can be set in the AUC tab.
 - Compute Confidence Interval (CI): check if you want to compute the CI. More details can be set in the CI tab.
 - Plot: check this box to create a plot. More details can be set in the Plot tab.
- Options
 - Percent: compute and report sensitivity and specificity in percent. If unchecked, fractions (between 0 and 1) are used instead.
 - Direction: the direction of the comparisons made to assess sensitivity and specificity. The comparison is made as "controls <> cases". "auto" determines the direction using the median.
- State
 - State variable: this is the variable defining the two groups that the Test variable discriminates.

- Level for controls: the value of State variable for the control observations.
- Level for cases: the value of State variable for the case observations.
- Output
 - Report: uncheck if you do not want the report to display.
 - Coordinates: how to report the coordinates of the ROC curve. "no" reports nothing, "local maximas" reports only thresholds corresponding to a local maxima of the ROC curve and "all" reports all the thresholds.
- Save ROC curve
 - Save as: the ROC curve will be saved under this name. You can later access it from the command window.

3.2 Smoothing tab



- Smooth: check if you want to perform smoothing.
- Smoothing method: the type of fit that will be employed to smooth the curve. "binormal" is a simple linear fit that will perform well in most cases. "density" will estimate a density distribution for cases and controls separately and deduce a smoothed ROC curve. With "fitdistr" you can specify a distribution for controls and cases.
- Density options: for "density" smoothing.
 - Bandwidth: the width of the window. "bcv", "ucv", "sj", "hb" and "nrd" are documented in TIBCO Spotfire S+ Language Reference as bandwidth.* functions. "<custom>" permits to specify a numeric value in the following field.
 - Numeric bandwidth: if "bandwidth" is "<custom>", a numeric value specifying the width of the window.

- Adjustment: the bandwidth will be multiplied by Adjustment. 1 will have no effect.
- Window: the kind of window. Note that all the windows are symmetric.
- Distribution options:
 - Distribution of controls, distribution of cases: density distributions for control and cases observations.

3.3 AUC tab

The screenshot shows the 'Build a ROC curve' dialog box with the 'AUC' tab selected. The 'General' tab is also visible. The 'Compute Area Under Curve (AUC)' checkbox is checked. The 'Partial AUC' checkbox is also checked. In the 'Partial AUC' group, the 'From' field is set to 100 and the 'To' field is set to 90. The 'Focus' group has 'specificity' selected with a radio button, and 'sensitivity' is unselected. The 'Correct partial AUC' checkbox is checked. At the bottom, there are buttons for 'OK', 'Cancel', 'Apply', '< >' (navigating between tabs), 'current' (the active tab), and 'Help'.

- Compute Area Under Curve (AUC): uncheck if you do not want to compute the AUC.
- Partial AUC: compute only a portion of the AUC, as defined in the Partial AUC group:
 - From, to: the bounds of the partial AUC. If percent was checked on the General tab, enter the numbers as percent (0-100), otherwise as proportion (0-1).
 - Focus: if the bounds are expressed in terms of sensitivity or specificity.
 - Correct partial AUC: correct the partial AUC in order to always have an area of 1 (or 100%) for a perfect partial discrimination and 0.5 (50%) for no discrimination, whatever the region specified. The formula by McClish (1989) is used.

3.4 CI tab

- Compute Confidence Interval (CI): check if you want to compute the CI.
- Confidence level: always between 0 and 1 (even if percent was checked on the General tab).
- Type of CI:

Note: modifications in these fields will be reflected in the General tab / output coordinates field and in the Plot tab / Plot CI, only if they have not been modified before.

 - Type of CI: "auc" to compute the CI of the AUC (as specified in the AUC tab), "thresholds" for the CI of thresholds (as specified in the Thresholds field just below), "se" to compute the CI of the sensitivity at fixed specificities, or "sp" to compute the CI of the specificity at fixed sensitivities.
 - Thresholds: if the type of CI is "thresholds", which thresholds must be tested? "all", "local maximas", or "<custom>" to define the threshold(s) in the Value field.
 - Value(s): if type of CI is "se" or "sp", or if "thresholds" is "<custom>", enter the values on which the CI will be assessed. For "se" or "sp", make sure to be consistent with the Percent checkbox in the General tab, or leave blank so the program will guess default values. You cannot leave the field blank for a custom threshold.

Note: you can enter several values with an S+ expression, for example `c(0.95, 0.9, 0.85)` or `seq(0, 1, 0.1)`.
- Confidence interval options:
 - Number of replicates: how many replicates for the CI? Higher numbers give a better approximation but take more time to compute

- Stratified: controls and cases are sampled separately, so that each replicate contains exactly the same number of them.
If unchecked, a non-stratified bootstrap is performed, and some replicates may not contain any case or control if the data is small or unbalanced. In such a case, a warning will be displayed and the resulting missing values will be ignored.

3.5 Plot tab

The screenshot shows the 'Build a ROC curve' dialog box with the 'Plot' tab selected. The dialog has five tabs: General, Smoothing, AUC, CI, and Plot. The 'Plot' tab contains the following settings:

- Plot:** ☒ Plot
- General:**
 - ☒ Identity line
 - ☐ Add to existing plot
 - ☒ Grid
 - Plot CI:
- Display Information:**
 - ☒ AUC
 - Coordinates:
 - Custom threshold(s):
- Polygons of AUC:**
 - ☐ Polygon of AUC
 - ☐ Polygon of maximal AUC
- Plot labels:**
 - Title:
 - X axis:
 - Y axis:

At the bottom, there are buttons for OK, Cancel, Apply, a navigation bar with arrows and 'current', and a Help button.

- Plot: check this box to create a plot.
- General:
 - Identity line: check to display the diagonal line showing the identity (no discrimination power)
 - Add to existing plot: if unchecked, a new plot will be created. Otherwise, the ROC curve will be added to an existing one. Don't check if there is no current plot, or if the current plot is not a ROC curve.
 - Grid: display a grid behind the ROC curve
 - Plot CI: "no": don't plot the confidence interval (only value if you didn't select confidence intervals in the General or CI tabs). "bars": display the intervals as bars. "shape": only if the CI is of type "se" or "sp": show a confidence shape rather than bars.
- Display Information
 - AUC: print the value of the AUC in the center of the plot.
 - Coordinates: display the thresholds with specificity and sensitivity for a selected set of ROC points. "no" inhibits the printing, "best" shows only the points corresponding to the maximal sum of sensitivity and specificity, "all" prints all the thresholds and "local maximas" shows a point with the numeric coordinates for each local maxima of the ROC

curve. If "custom", fill-in the following field with the threshold(s) you want to display on the plot. Only "no" and "best" are available for smoothed ROC curves, and only sensitivity and specificity will be printed.

- Custom threshold(s): if "custom" is selected in "Coordinates", this field selects the threshold(s) that must be plotted. You can enter several thresholds with the following syntax: `c(5, 8.5)` to display the thresholds at 5 and 8.5.
- Polygons of AUC:
 - Polygon of AUC: highlight the AUC.
 - Polygon of maximal AUC: highlight the maximal region the AUC could take.
- Plot labels:
 - Title: the title above the graph, by default the name of the Test variable
 - X axis, Y axis: the labels of the X and Y axis.

4. Comparing ROC curves

The "ROC curves comparison" dialog serves to make a statistical comparison of 2 (partial) area under the ROC curve.

4.1 General tab

The screenshot shows the 'Compare two ROC curves' dialog box with the 'General' tab selected. The dialog has five tabs: General, Smoothing, AUC, Bootstrap, and Plot. The 'General' tab contains the following fields and options:

- Data Selection:**
 - Data Set: SubArachnoidH
 - ☒ Remove NAs
 - Test variable 1: wfns
 - Test variable 2: s100b
- Options:**
 - Method: auto
 - ☒ Percent
 - Direction: auto
- State:**
 - State variable: outcome
 - Level for controls: Good
 - Level for cases: Poor
- Output:**
 - ☒ Report
- Save ROC curve:**
 - Save as: last.ROCtest

At the bottom of the dialog are buttons for OK, Cancel, Apply, a navigation set (left arrow, right arrow, and 'current'), and a Help button.

- Data Selection
 - Data Set: the name of the data frame to use.
 - Remove NAs: if one of the Test or State variable contains missing values, remove the whole row of the data.

- Test variable1, Test variable 2: the names of the two test (continuous) variables defining the two ROC curves to compare.
- Options
 - Method: "delong" (only for comparing complete AUCs) or "bootstrap". With "auto", the appropriate method will be determined automatically.
 - Percent: compute and report sensitivity and specificity in percent. If unchecked, fractions (between 0 and 1) are used instead.
 - Direction: the direction of the comparisons made to assess sensitivity and specificity. The comparison is made as "controls <> cases". "auto" determines the direction using the median.
- State
 - State variable: this is the variable defining the two groups that the Test variables discriminates.
 - Level for controls: the value of State variable in the control observations.
 - Level for cases: the value of State variable in the case observations.
- Output
 - Report: uncheck if you do not want the report to display.
- Save ROC curve
 - Save as: the comparison object (htest) will be saved under this name. You can later access it from the command window.

4.2 Smoothing tab

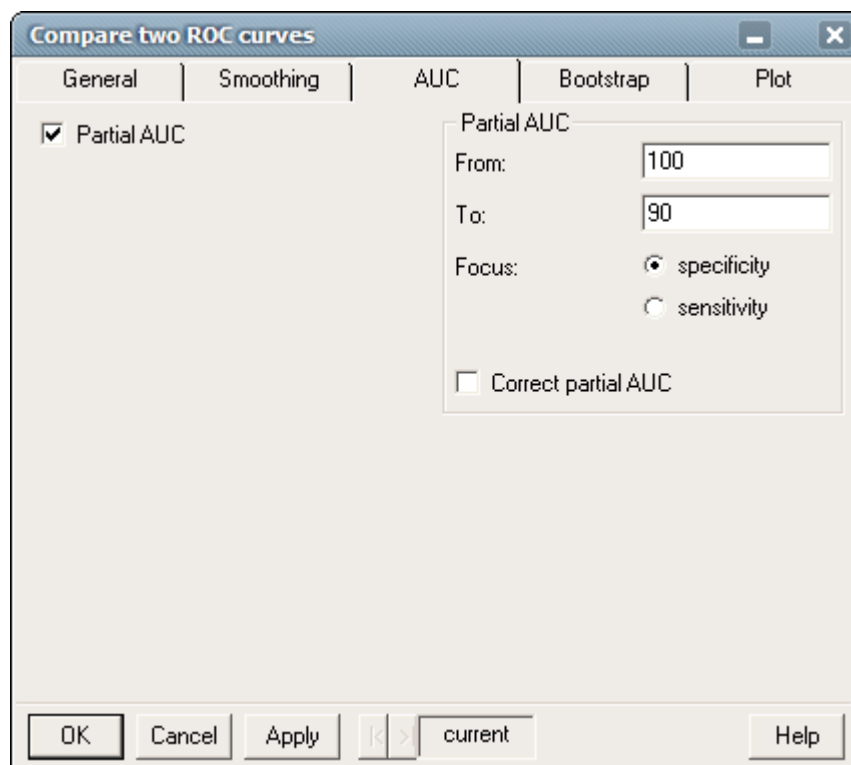
The screenshot shows the 'Compare two ROC curves' dialog box with the 'Smoothing' tab selected. The dialog is divided into two columns for 'Smooth ROC 1' and 'Smooth ROC 2'. Both columns have the 'Smooth' checkbox checked. The 'Smoothing method' is set to 'density' for both. Under 'Density options', 'Bandwidth' is set to '<numeric>', 'Numeric bandwidth' is empty, 'Adjustment' is set to '1', and 'Window' is set to 'gaussian'. Under 'Distributions options', both 'Distribution of controls' and 'Distribution of cases' are set to 'normal'. The bottom of the dialog contains buttons for 'OK', 'Cancel', 'Apply', '< >' (disabled), 'current', and 'Help'.

- Smooth ROC 1, Smooth ROC 2: smooth the ROC curves corresponding to Test variable1 and Test variable 2.
- Smoothing method: the type of fit that will be employed to smooth the curve. "binormal" is a simple linear fit that will perform well in most

cases. "density" will estimate a density distribution for cases and controls separately and deduce a smoothed ROC curve. With "fitdistr" you can specify a distribution for controls and cases.

- Density options: for "density" smoothing.
 - Bandwidth: the width of the window. "bcv", "ucv", "sj", "hb" and "nrd" are documented in TIBCO Spotfire S+ Language Reference as bandwidth.* functions. "<custom>" permits to specify a numeric value in the following field.
 - Numeric bandwidth: if "bandwidth" is "<custom>", a numeric value specifying the width of the window.
 - Adjustment: the bandwidth will be multiplied by Adjustment. 1 will have no effect.
 - Window: the kind of window. Note that all the windows are symmetric.
- Distribution options:
- Distribution of controls, distribution of cases: density distributions for control and cases observations.


4.3 AUC tab



- Partial AUC: compute only a portion of the AUC, as defined in the Partial AUC group (not available if "method" (in "General" tab) is "delong"):
 - From, To: the bounds of the AUC. If "percent" was checked on the "General" tab, enter the numbers as percent (0-100), otherwise as proportion (0-1).
 - Focus: if the bounds are expressed in terms of sensitivity or specificity.


- Correct partial AUC: correct the partial AUC in order to always have an area of 1 (or 100%) for a perfect partial discrimination and 0.5 (50%) for no discrimination, whatever the region specified. The formula by McClish (1989) is used.

4.4 Bootstrap tab



- Confidence interval options (not available if "method" (in "General" tab) is "delong"):
 - Number of replicates: how many replicates for the CI? Higher numbers give a better approximation but take more time to compute
 - Stratified: controls and cases are sampled separately, so that each replicate contains exactly the same number of them. If unchecked, a non-stratified bootstrap is performed, and some replicates may not contain any case or control if the data is small or unbalanced. In such a case, a warning will be displayed and the resulting missing values will be ignored.

4.5 Plot tab



- Plot: check this box to plot the two ROC curves compared.
- General:
 - Identity line: check to display the diagonal line showing the identity (no discrimination power)
 - Grid: display a grid behind the ROC curves
- Legend: the legend that will be displayed in the plot.
 - Legend of ROC 1, Legend of ROC 2: legend annotation.
 - Color of ROC 1, Color of ROC 2: the colors of the two ROC curves in the plot, corresponding to the legend.
- Polygon of maximal AUC: highlight the maximal region the AUC could take.
- Plot labels:
 - Title: the title above the graph.
 - X axis, Y axis: the labels of the X and Y axis.

5. Algorithms

5.1 Area Under the Curve

The AUC is computed with the trapezoidal rule as described by Fawcett (2006). The partial AUC has been theoretized by McClish (1989). It is computed with the trapezoidal rule with interpolation if necessary. The partial AUC correction is done as described by McClish (formula 6).

5.2 Confidence intervals

The CI are computed with bootstrap. In stratified bootstrap, control and case observations are sampled independently, otherwise all the observations are sampled at the same time.

After "Number of replicates" repetitions, the quantiles of the resulting statistics are computed at $(1 - \text{"confidence level"})/2$, 0.5 and $1 - (1 - \text{"confidence level"})/2$.

5.3 ROC curve comparison

5.3.1 Bootstrap

With the "bootstrap" method, the processing is done as follow:

1. "Number of replicates" bootstrap replicates are drawn from the data. If "stratified" is checked, each replicate contains exactly the same number of controls and cases than the original sample, otherwise the numbers can vary.
2. for each bootstrap replicate, the AUC of the two ROC curves are computed and the difference is stored.
3. The following formula is used:

$$D = (AUC1 - AUC2) / s$$

where s is the standard deviation of the bootstrap differences and AUC1 and AUC2 the AUC of the two (original) ROC curves.

4. D is then compared to the normal distribution to compute the p-value.

5.3.2 DeLong

With the "delong" method, the processing is done as described in DeLong *et al.* (1988). Only comparison of two ROC curves is implemented.

6. References

James Carpenter and John Bithell (2000) Bootstrap condence intervals: when, which, what? A practical guide for medical statisticians. *Statistics in Medicine* 19, 1141-1164.

Elisabeth R. DeLong, David M. DeLong and Daniel L. Clarke-Pearson (1988) Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics* 44, 837-845.

Tom Fawcett (2006) An introduction to ROC analysis. *Pattern Recognition Letters* 27, 861-874. DOI: 10.1016/j.patrec.2005.10.010.

James A. Hanley and Barbara J. McNeil (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* 143, 29-36.

Donna Katzman McClish (1989) Analyzing a Portion of the ROC Curve.
Medical Decision Making 9(3), 190-195. DOI:
10.1177/0272989X8900900307

Xavier Robin, Natacha Turck, Jean-Charles Sanchez and Markus Müller
(2009) Combination of protein biomarkers. *useR! 2009*, Rennes.