

5.3 ASTSA



ASTSA is an interactive, menu driven time series package. Routines are selected from the menus. Some information is typed into the input boxes. Each routine produces some output, text and/or graphs. Text output appears in the top window, graphs appear in the lower window or full screen. These windows can be moved and resized. The windows can be left opened and the user can return to the main ASTSA menu to run another routine. The results from many routines can be left opened for comparison.

Each graph window, see example below, shares certain features. Moving the mouse over the plot gives x , y coordinates in the lower left hand corner of the window. X is usually the time index t . Y is the value of the series. Portions of a graph may be zoomed to allow for more detail. To zoom, move the mouse cursor over the graph and choose a corner of the region to be zoomed. Hold down the left button and move the mouse until the rectangle encompasses the desired region. Release the button. Move the mouse cursor inside this region and click the left button to zoom this region. Move the mouse cursor outside this region and click the left button to cancel the zoom. This process may be repeated. Double click on the graph to return to the original view. There is no x , y readout when the zoom rectangle is sized.

In windows with multiple graph plots, each graph may be zoomed individually. Double click on the legend of the desired graph to remove the other graphs from view and resize the selected graph to full size. The All plots menu heading will darken. Choose the All plots menu option to return to the multiple graph view. There may be times when one of the graphs is not needed. Press the right mouse button over this graph to remove it from view. The remaining graphs will enlarge to fill the space and the All plots menu heading will darken. Choose the All plots menu option to return to the multiple graph view.

Unless mentioned otherwise, all graph windows share the following menu items: File, Grid, Annotate, Quit. Under the File menu option, we find four submenus, Print, Page setup, Printer setup, and Store.

Print sends a copy of the graph to the printer.

Page setup allows the user to position the graph on the paper. Page setup also allows the selection of font and font sizes for the printed graph.

Printer setup allows the selection of printers, if more than one printer is available on the user's machine, and the print resolution. Usually, the printers default settings work well and this option is rarely needed.

Store saves a copy of the graphs to disk. These saved files can be found in the graphs sub directory under the ASTSA directory. Saved graphs can be viewed using the view graphs menu option described later.

Grid toggles the grid lines. Choose grid once to remove the grid lines. Choose grid again to draw the grid lines.

Annotate includes the following submenus, Title, X axis legend, Y axis legend, Add text, Edit text, X axis scale, and Y axis scale. Some routines include other options under annotate.

Title allows the user to change the title of the graphs.

X axis legend and Y axis legend allows the user to customize the axes descriptions.

Add text allows the user add text to the graph. When this menu is selected, a text box appears in the middle of the graph. Press and hold the left mouse button to position the

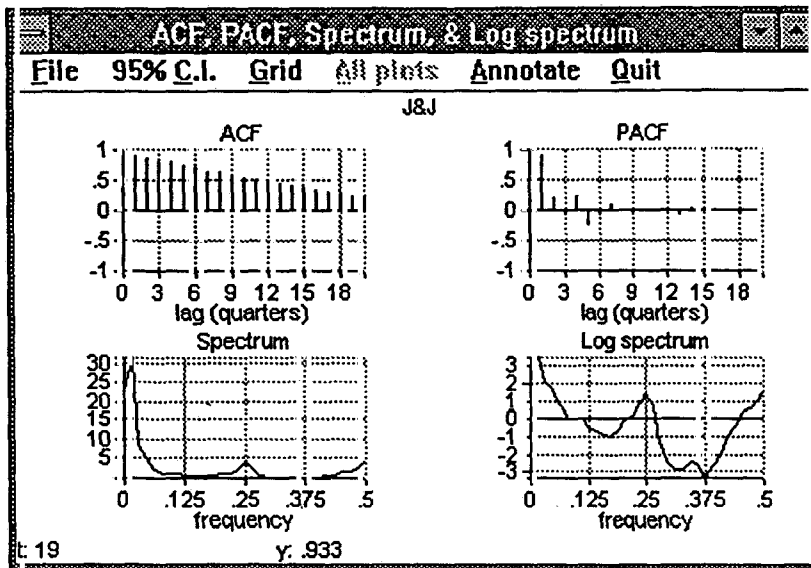
text box at the desired location. Enter your text and press return when finished. The text now appears on the graph and the text box returns to the center of the graph. Choose the add text menu again to remove this text box.

Use edit text to edit text previously added to the graph via the add text option.

X axis scale changes the minimum and maximum x values (usually the time index) as well as the x increments. This option can be used to zoom the graph with the keyboard.

Y axis scale does the same for the y axis scaling.

Quit closes the windows and returns the user to the ASTSA main menu.



An example of a graph window showing the menu items File, Grid, Annotate, and Quit.

Some routines produce text output. The text output window displays the results used to generate the graph and other details relevant to the routine. Scroll through the output using the mouse and the scroll bars. Click on the scroll bar arrow keys to scroll one line of text. Drag the gray box in the scroll bar for rapid scrolling of text. The text scrolls as the mouse drags the gray box. ASTSA appends all output to the log file which maybe reviewed at a later time. The text window has one or two menu items, Print and sometimes New page. Print sends a copy of the results to the printer. A smaller portion of the text can be selected by holding the left mouse button and dragging the mouse over the text. This selected text is highlighted and the portion highlighted is controlled by dragging the mouse. Click anywhere in the text to remove the highlight. If the print menu item is selected when some portion of the results is highlighted, a print window appears allowing the user to choose whether to print the full output or just the selected portion. When a very large amount of output is generated, the output must be broken into blocks or pages. When this happens, the New page menu item appears. Choose the New page menu to view the list of pages. Click on a number to select another page of output.

Output					
Print		New page			
J&J					
ACF					
zero lag autocovariance = 18.576					
lag	ACF	lag	ACF	lag	ACF
1	.9251	2	.8883	3	.8328
5	.7638	6	.7176	7	.6750
9	.6083	10	.5642	11	.5256

An example of the text output. Here the output is long and the New page menu item is included.


Most routines in ASTSA require more input from the user than the standard windows input box can supply. Thus ASTSA uses its own input boxes. Shown below. The lines may be navigated using the cursor keys and the page up/down keys. Press home to position the cursor at the beginning of the current line. Press ctrl-home to move to the beginning of the first line. Press end to move to the end of the current line. Press ctrl-end to move to the end of the last line. For routines requiring large amounts of input, a scroll bar appears. The mouse may be used with the scroll bar. Click on the arrows to scroll one line, click on the scroll bar to scroll many lines. Return acts similar to the down arrow but positions the cursor at the beginning of the next line. When return is pressed on the last line, the OK button lights up. Press return again to select OK. Note that the series refers to the series ID number found in the main menu. The series name is provided for descriptive purposes only.

Spectrum	
series: 1	J&J
T' (128):	128
number of estimates:	32
smoothing:	3
cosine taper (y/n):	n
<input type="button" value="OK"/> <input type="button" value="Cancel"/>	

An example of the input box for the spectrum routine.

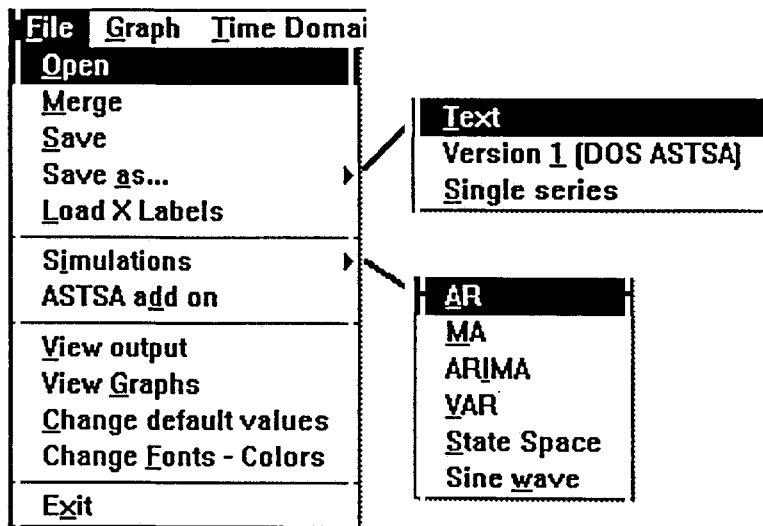
Organization of the routine descriptions parallels the layout of the menus. We begin with the main ASTSA window.

The main ASTSA window and menu showing the list of series.



FILE

The following routines are found under the file menu option. Including the six simulations found under the simulation menu item.



Open

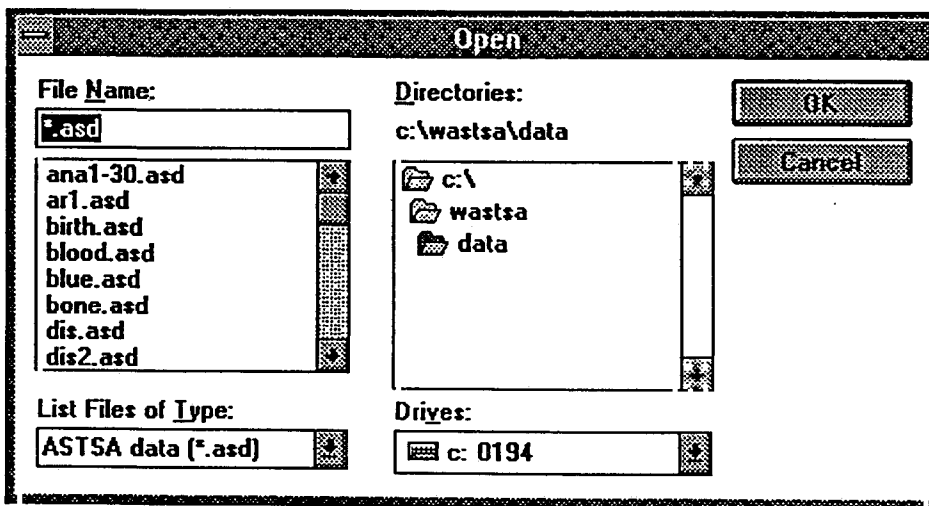
Open inputs data into ASTSA from either a text data file or native ASTSA .ASD file. Note that the open routine will delete all series currently in ASTSA and replace them with the new series in the data file. Open defaults to the .ASD file format. Other file extensions can be chosen in the file window shown below.

The file window lists the .ASD files found in the data sub directory. Double click on the file name to open or single click on the file name and then click on the OK button. Click on cancel to return to the main menu. Files can also be selected by typing the file name in the file name box and then click on OK. Directories can be navigated by double clicking on the file folder icons in the directories box. Drives can be changed by selecting a different drive letter from the drive box.

The native ASTSA file format requires no input from the user. All information about the data is stored in the file along with the actual data. Retrieving data from a text file requires additional information from the user. Series names and other information must be supplied by the user. Series names must not begin with numbers. Names left blank are supplied with letter names. The file window appears first listing all files with the .DAT extension. Once a file is selected, the input box appears.

Input: The input box for Import is of varying size. Initially prompts the user for:
No. of series: Number N of series in the data file.
Time units: Time units of the data (e.g. months, days).

Once the number of series is chosen, the input box expands and prompts for:
No. of series: Number N of series in the data file.
Series 1: Name of first series. Must not start with a number.
:
Series N: Name of last series. Must not start with a number.
Time units: Time units of the data (e.g. months, days).



The file window.

Merge

Similar to open but existing series in ASTSA are not erased. This option allows data from different files to be combined in one ASTSA session. The user can merge another .ASD file into the existing ASTSA session, or merge data from a text file into the existing ASTSA session. Merging a .ASD file requires no further input from the user. Merging text requires more information from the user. After selecting the text file, the input box appears. This input box is the same as the input box described above in the open menu option.

Save

Saves the current data, series names and options into an ASTSA .ASD binary file. Save takes a snapshot of the current working conditions. The file window appears first listing all files with the .ASD extension. Select a file or type a new file name in the file box. Click on the OK button when a file is selected. If the file exists, an overwrite message box will appear. Choose OK to replace the file with the current data. Choose cancel to return to the file window.

Save as ...

Saves the current data into an text file, a version 1 file, or saves a single series. In all cases, the file window appears. Select a file or type a new file name in the file box. Click on the OK button when a file is selected. If the file exists, an overwrite message box will appear. Choose OK to replace the file with the current data. Choose cancel to return to the file window.

Saving with the text option saves multiple series end-to-end as one long column in the file. All series lengths must be equal and thus shorter series are padded with zeros. Text files default to the .DAT extension.

Version 1 files also default to the .ASD extension. ASTSA can distinguish the difference between version 1 and version 2 files.

To save a single series, first click on the series in the main list. Then the file window appears. If no series is selected, an error message appears.

Load X labels

ASTSA displays the time index as numerical values. Often, more meaningful labels for the time index should be used, such as month/year should be used. Unfortunately, there is no way to automate such time labels in ASTSA. However, the labels can be stored in a text file which ASTSA can read. Each label must be less than 20 characters in length. One label for each time index is required. However, if forecasting is desired, future values will be generated by ASTSA. A text file may contain more labels than actual data points so that these future values can have meaningful labels.

Simulations

Allows the user to generate sine waves and simulate AR, MA, ARIMA, VAR, and state space models. ASTSA generates 50 points longer than the length required. The series start at 1 and the first part of the simulated series is discarded, only the last t values are kept. Each simulation requires some information from the user. Note that no prior variance is used for state-space simulations. Sine waves are not simulated but they generate contrived series and are thus group with the simulations. All the input boxes for simulations are of varying size. Once the orders of the models are selected, the input boxes expand. Note that some parameter choices lead to non stationary and unstable models. In such cases, the simulation routine may crash. If the simulation routine crashes with an overflow error, this usually indicates that the model is unstable and a different set of parameters should be used.

AR input: The input box for AR simulations prompts the user for:

- Order: Order p of the autoregression.
- Length t : Length of the simulated series.
- Constant: Value of the constant term.
- Noise var: Variance of the white noise.
- $\phi(1)$: Parameter value.
- :
- $\phi(p)$: Parameter value.

MA input: The input box for MA simulations prompts the user for:

- Order: Order q of the moving average.
- Length t : Length of the simulated series.
- Noise var: Variance of the white noise.

theta(1): Parameter value.

:

theta(q): Parameter value.

ARIMA input: The input box for ARIMA simulations prompts the user for:

p: Order p of the autoregression.

d: Order d of the difference.

q: Order q of the moving average.

P: Order P of the seasonal autoregression.

D: Order D of the seasonal difference.

Q: Order Q of the seasonal moving average..

Season: Length of the seasonal component.

Length t: Length of the simulated series.

Noise var: Variance of the white noise.

phi(1): AR parameter value. Appears if $p > 0$.

:

phi(p): AR parameter value. Appears if $p > 0$.

theta(1): MA parameter value. Appears if $p > 0$.

:

theta(q): MA parameter value. Appears if $p > 0$.

PHI(1): Seasonal AR parameter value. Appears if $p > 0$.

:

PHI(P): Seasonal AR parameter value. Appears if $p > 0$.

THETA(1): Seasonal MA parameter value. Appears if $p > 0$.

:

THETA(Q): Seasonal MA parameter value. Appears if $p > 0$.

VAR input: The input box for VAR simulations prompts the user for:

Number of series

q:

Order of VAR

p:

Length of simulated series

t:

Constant vector of dim q

const 1: Value of the constant term for the first series.

:

const q: Value of the constant term for the last series.

Phi 1 of dim q x p

row 1: First row of the phi matrix. P columns of values separated by spaces.

:

row q: Last row of the phi matrix. P columns of values separated by spaces.

:

Phi p of dim q x p

row 1: First row of the phi matrix. P columns of values separated by spaces.

:

row q: Last row of the phi matrix. P columns of values separated by spaces.

Error covariance matrix of dim $q \times q$
row 1: First row of the phi matrix. Q columns of values separated by spaces.
:
row q : Last row of the phi matrix. Q columns of values separated by spaces.

State space input: The input box for state space simulations prompts the user for:
Dimension of observation vector q (no. of series)
 q :
Dimension of state vector p
 p :
Length of simulated series
 t :
Initial X_0 values of dim p
 x_0 1: First component of the x_0 vector. Only one value.
:
 x_0 p : Last component of the x_0 vector. Only one value.

Transition matrix PHI of dim $p \times p$
row 1: First row of the phi matrix. P columns of values separated by spaces.
:
row p : Last row of the phi matrix. P columns of values separated by spaces.

State covariance matrix Q of dim $p \times p$
row 1: First row of the Q matrix. P columns of values separated by spaces.
:
row p : Last row of the Q matrix. P columns of values separated by spaces.

Measurement matrix A of dim $q \times p$
row 1: First row of the A matrix. P columns of values separated by spaces.
:
row q : Last row of the A matrix. P columns of values separated by spaces.

Observation covariance matrix R of dim $q \times q$
row 1: First row of the R matrix. q columns of values separated by spaces.
:
row q : Last row of the R matrix. q columns of values separated by spaces.

Sine input: Although not a true simulation. A noise component can be added to the sine wave and thus sine waves are found together with the simulations. Sine waves are generated by $x(t)=a\text{SIN}(2\pi(t+s)/p)$. The input box for sine simulations prompts the user for:
Period: Period length p in time points for the sine waves
Amplitude: Amplitude a of the sine wave.
Shift: Shift s in time points for the sine wave. This option can effectively generate cosine waves, by setting s to $1/4$ period length.
Length t : Length of the generated series.
Noise var: Variance of the white noise added to the sine wave.

ASTSA add on

This as an advanced feature for those who want to write there own routines that access ASTSA data.

View output

ASTSA maintains a running log file of all results. Most of the routines produce text output or results. This text is displayed when the routine is finished. It is also appended to the log file. This log file may be viewed and or printed at any time. All or a portion of the file may be printed. Portions of the output may be highlighted by dragging the mouse and printed. There are several menu options for viewing.

Print: Prints the entire log file, or the highlighted portion.

Copy: Copies a selected portion of text to the windows clipboard. Once in the window, the text can be copied into other windows programs.

Save: Renames the log file and starts a new log file. The file window appears allowing the user to select a file name.

Delete: Erases the log file and starts a new log file.

Quit: Closes the viewing window and returns the user to the main ASTSA menu.

Page: If the log file is long, a page option may appear. Page allows the user to select another block of the log file to review.

View Graphs

Graphs produced by ASTSA may be saved to disk with the store option. These saved files can be found in the GRAPHS sub directory under the ASTSA directory. Saved graphs have the Windows .BMP bit map extension and are numbered sequentially GRPH0001.BMP, GRPH0002.BMP, etc. An information file stores the title of the graph as well as the graph file number. The list on the right is the list of stored graphs. Click on a name to view the graph. The graph window can be sized to view the whole graph. Three menu options.

Print: Sends a copy of the selected graph to the printer.

Delete: Deletes the selected graph from the list and disk.

Quit: Closes the viewing window and returns to main ASTSA menu.

Change default values

Allows the user to change some of ASTSA's default values for some of the time domain and frequency domain routines. These values appear in the input boxes.

Input: The change default values box prompts the user for:

Maximum lag (≥ 1): Maximum lag used in ACF, PACF, CCF, and variogram routines.

T: Extended power of 2 series length used in frequency domain routines. The numbers in parentheses are the next three valid powers of 2. T must be \geq actual length of the series.

Number of spectral estimates: Number of spectral estimates displayed. Must be a power of 2 smaller than T. Used by the frequency domain routines.

Smoothing: Amount of smoothing for the above number of spectral estimates. Must be an odd number.

Cosine taper (y/n): Determines whether to perform a cosine taper on the data. Used by the frequency domain.

Grid: Determines if the grid is initially drawn for each plot. May be toggled in the plot window. Choose 0 for no grid, 1 to draw the grid.

Time units: Each time series point is collected at regular increments, this increment is the units.

Selector (AIC, AICc, BIC, FPEL, AICL): Choice of model selection criterion for the modeling routines such as ARIMA.

Change Fonts - Colors

Invokes the ASTSA font program and is the same as double clicking on the ASTSA fonts icon. Allows the user to change the fonts and text colors for ASTSA. This is a separate program from ASTSA and is described later in section 5.5.

Exit

If nothing is in the log file, prompts with a quit ASTSA message. Choose yes to quit, no to return to the main menu. If something is in the log file, prompts with a save log file message. Choose yes to save the log file and quit, choose no to delete the log file and quit, choose cancel to return to the main menu. If the log file is saved, the file window appears. Select a file or type file name in the file box. Choose OK. If the file exists, an overwrite message appears. Double click the control box in the upper left hand corner of the main ASTSA menu window to quit ASTSA without deleting the list of series.

GRAPHS

Graphs include descriptive routines as well as the simple graphing routines. Covariance matrix and residual tests are descriptive in nature and fall under this category. The last option, All 4 ACF, PACF, spectrum, log spectrum falls under graphs since it belongs to both the time domain and the frequency domain.

<u>G</u> raph	<u>T</u> ime Domain	<u>F</u> requency Domain
Graph series		
Plot 2 series		
Multiple plots		
Plot All		
2-D scatter plot		
Multiple 2-D plots		
3-D scatter plot		
Coded 2-D scatter plot		
Stacked plot		
Moving statistics		
Histogram		
Covariance matrix		
Residual tests		
All 4 ACF, PACF, Spectrum, Log spectrum		

The following routines are found under the graph option.

Graph series.

Displays a plot of the series selected. Graph series can be invoked in one of two ways. First, use the graph series menu option. Second, double click on a series in the main list. Double clicking bypasses the input box.

Input: The input box for graph series prompts the user for:

Series: ID number or full name of the series.

Graph: Displays a plot of the selected series. The graph window has two additional submenu options under annotate.

Markers: Draws solid circles around each data point. Click once to add the markers. Click again to remove the markers.

Season: Draws vertical lines for seasonal data. An input box appears which prompts the user for:

Season: Length of the seasonal component.

Shift: Starting point for the seasonal component.

Plot 2 series

Displays a plot of both series selected. The series can have individual or common scaling on the Y-axis. The first series is the double thick line. The second series is the thin line.

Input: The input box for plot 2 series prompts the user for:

Series1: ID number or full name of first series.

Series2: ID number or full name of second series.

Graph: Moving the mouse over the plot gives x, y coordinates in both lower corners of the window. Series one on the left and series two on the right. The graph window has one additional menu option.

Common scale: Toggles whether the series share a common Y-axis. Does not work when portions of the plots are zoomed.

Multiple plots

Displays a 2x2 grid of graphs of each series. If there are more series than four series, a scroll bar appears. The graphs can be scrolled. These plots are not zoomable. There is no input box required for this routine. Multiple plots contains one additional menu option, positions.

Positions: An input box appears allowing the user to change the number of rows and columns in the display grid.

Plot All

Displays all series scaled such that each series maximum touches the top and each series minimum touches the bottom of the graph.

Graph: No x, y readout or zoom feature and no grid line option.

2-D scatter plot

Displays 20 scatter plots of series1(t) vs. series2(t-lag) at lags: lag, lag+1, lag+2, ... lag+19. The lag appears in the upper left of each scatter plot. A single detailed scatter plot including correlation can also be displayed.

Input: The input box for 2-D prompts the user for:

Series1: X axis. ID number or full name of first series.

Series2: Y axis. ID number or full name of second series.

Lag: Lag of second series.

Full or Detail view (f/d): Choose f to see the full 20 lagged scatter plots first. Choose d to see the detail scatter plot first.

Graph: On the 20 plot display, double click on one of the squares for a detailed plot including correlation. Double click again to return to the 20 plot display. On the detail display, moving the mouse over the graph gives x, y coordinates in the lower left hand corner of the window. X is the value of series1 and y is the value of series2. No zoom feature. The graph window has one additional menu option.

Dot size: Changes each plotted point's dot size. Choose from small, medium, or large. The current size is check marked.

Multiple 2-D plots

Displays scatter plots of all series vs. each other at a given lag. Scroll bars appear when the number of series is greater than five. A single detailed scatter plot including correlation can also be displayed. If the data set contains more than five series, scroll bars appear in the many graph view. Use the scroll bars to view the other series.

Input: The input box for multiple 2-D prompts the user for:

Lag: Time lag.

Graph: On the multiple plot display, double click on one of the squares for a detailed plot including correlation. Double click again to return to the multiple plot display. On the detail display, moving the mouse over the graph gives x, y coordinates in the lower left hand corner of the window. No zoom feature. The graph window has one additional menu option.

Dot size: Changes each plotted point's dot size. Choose from small, medium, or large. The current size is check marked.

Coded 2-D scatter plot

Similar to the usual 2-D scatter plot routine with the addition that a third series of codes is required. The codes determine the shape of the marker for the data. A code of 1 displays circles, 2 displays squares, and 3 displays triangles. This routine is usually used with discriminate analysis.

Input: The input box for 2-D prompts the user for:

Series1: X axis. ID number or full name of first series.

Series2: Y axis. ID number or full name of second series.

Code series: ID number or full name of code series.

Lag: Lag of second series.

Graph: On the 2D plot display, double click on one of the squares for a detailed plot including correlation. Double click again to return to the 2D plot display. On the detail display, moving the mouse over the graph gives x, y coordinates in the lower left hand corner of the window. X is the value of series1 and y is the value of series2. No zoom feature. The graph window has one additional menu option.

Dot size: Changes each plotted point's dot size. Choose from small, medium, or large.
The current size is check marked.

3-D scatter plot

Displays a rotating scatter plot of series1(t) vs. series2(t-lag1) vs. series3(t-lag2). Initially spins to the right at medium speed.

Input: The input box for 3-D scatter plot prompts the user for:

Series1: X axis. ID number or full name of first series.
Series2: Y axis. ID number or full name of second series.
Lag: Time lag of second series. Must be positive.
Series3: Z axis. ID number or full name of third series.
Lag: Time lag of third series. Must be positive.

Graph: The rotation speed is controlled by the scroll bar. Use the mouse to click on the scroll bar arrows or slide the indicator. The graph window has the following menu options.

Up: Starts the cloud rotating upwards.
Down: Starts the cloud rotating downwards.
Left: Starts the cloud rotating to the left.
Right: Starts the cloud rotating to the right.
Print: Freezes and sends a copy of the graph to the printer.
Help: Displays the 3D plot portion of the help file.
Quit: Closes the windows and returns to main ASTSA menu.

Note that the menu heading for the current direction is gray.

Stacked plot

Stacks the series in runs of w, the window size. Useful for seasonal data.

Input: The input box for stacked plot prompts the user for:

Series: ID number or full name of the series.
Window: Length of the runs in the series.

Graph: Displays a surface plot of the stacked series. The runs of length w is labeled w. The stack of runs is labeled s. The graph window has one additional menu option.

Surface: Determines the surface type, lines or a grid. Four options. The current choice is check marked.
Line s: Draws parallel lines in one direction.
Line w: Draws parallel lines in the other direction.

Grid: Draws a grid over the surface. Lines in both directions.

Hidden Line Removal: Toggles the hidden line removal. Choose once to turn this feature on, choose again to turn this feature off. Gives the illusion that the surface is solid and not transparent by removing the parts of the graph that should not be visible. Sometimes the hidden line removal fails in which case the parts of the graph are missing. Choose this option again to turn off the hidden line routine. Try a different angle and try removing the lines again.

View: Determines which quadrant is displayed. Ninety degree rotations. Four options, quad 1, quad 2, quad 3, quad 4. The current quadrant is check marked.

Moving statistics

Computes a running centered moving average and a running centered moving standard deviation over the selected series. This gives a visual check as to whether the mean and variance remain constant over time. The amount of smoothing (data to compute the moving average and standard deviation) depends on the window. Data at the ends equal to about half the window are deleted.

Input: The input box for moving statistics prompts the user for:

Series: ID number or full name.

Window: Smoothing size for moving average and standard deviation.

Graph: Displays the raw series in the top graph. A running moving average in the second graph. And a running moving standard deviation in the bottom graph. The graph window has one additional menu option.

All plots: Restores all three plots.

Histogram

Draws a sixteen bin histogram of the series and superimposes a normal curve. Also calculates the mean and standard deviation.

Input: The input box for histogram prompts the user for:

Series: ID number or full name.

Graph: No zoom feature.

Covariance matrix

Although not a graph, this routine computes the mean vector, covariance matrix, eigenvalues for the covariance matrix, inverse of the covariance matrix, correlation matrix, eigenvalues for the correlation matrix, Hotelling's T-square, determinant of the covariance matrix, and the AICc based on a constant only model.

Input: The input box for covariance is of varying size. Initially prompts the user for:

No. of series: Number of series included in the covariance matrix.

Once the number of series is chosen, the input box expands and prompts for:

No. of series: Number of series included in the covariance matrix.

Series1: ID number or full name of the first series.

:

SeriesN: ID number or full name of the last series.

Output: Displays the above results. These results are appended to the log file which maybe reviewed at a later time. The results window has several menu options.

Print: Sends a copy of the results to the printer.

Help: Displays covariance matrix portion of the help file.

Quit: Closes the window and returns to main ASTSA menu.

Residual tests

Examines the cumulative spectrum, normality QQ plot, fluctuations, outliers, Box-Pierce test for serial correlation at lags 1 and 20, and histogram of the selected series. The cumulative spectrum tests for flat spectrum and constant variance. The Box-Pierce tests for serial correlation at both lag 1 and lag 20. A white noise should not have too many fluctuations, changes from upwards trends to downward trends or visa-versa, or too few fluctuations. This tests counts the number of fluctuation or trend changes. For each of the T points, a z-score is calculated. The largest absolute z-scores is compared to the largest absolute z-score from T standard normals. The QQ plot measures the correlation between the ordered data and standard normal percentiles. The higher the correlation, the more closely the data follows a normal distribution.

Input: The input box for residual tests prompts the user for:

Series: ID number or full name.

Output: Displays length t of the series. The maximum absolute difference for the cumulative spectrum. Its p-value as either n.s, <.05, or <.01 with a comment on the result. The chi squared at lags 1 and 20 for the Box-Pierce. Their p-values as either n.s, <.05, or <.01 with a comment on the result. The z-score for the fluctuation test. Its p-value as either n.s, <.05, or <.01 with a comment on the result. The maximum absolute z-score for the outlier check. Its p-value as either n.s, <.05, or <.01 with a comment on the result. The correlation in the QQ plot. Its p-value as either n.s, <.05, or <.01 with a comment on the result. White noise should be non significant for all tests except the outlier test.

Graph: Displays the cumulative spectrum plot, the QQ plot, and the histogram with superimposed normal curve. No zoom feature. The graph window has two additional menu options.

All Plots: Displays all three graphs.

Dot size: Changes each plotted point's dot size in QQ plot. Choose from small, medium, or large. The current size is check marked.

Small

Medium

Large

All 4 ACF, PACF, spectrum, log spectrum

Computes the autocorrelation (ACF), and partial autocorrelation (PACF) functions for lags 0 to m. Computes spectrum and log spectrum for the chosen series.

Input: The input box for all 4 prompts the user for:

Series: ID number or full name.

Max lag: Maximum lag m for ACF, PACF. Calculates from 0 to m.

T': Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the results of ACF with zero lag autocovariance and 95% significance, PACF with 95% significance, spectrum with high and low confidence band multipliers, and log spectrum with high and low confidence band adders.

Graph: Graphs of the ACF, PACF, spectrum, and log spectrum are displayed in the bottom window. The graph window has two additional menu options.

95% sig: Toggles 95% significance/confidence lines on the graphs. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance. Significance lines for ACF, PACF. Confidence bands for spectrum, log spectrum.

All Plots: Restores all four graphs.

Time Domain

Includes both diagnostic and descriptive routines in the time domain. Also includes the models.

Time Domain	Frequency
<u>A</u> CF	
<u>P</u> ACF	
<u>V</u> ariogram	
<u>C</u> CF	
<u>A</u> utoregression	
<u>A</u> R <u>I</u> MA	
<u>A</u> R <u>I</u> MA search	
<u>M</u> ultiple regression	
<u>S</u> tate-Space	
<u>V</u> ector autoregression	
<u>M</u> ultivariate regression	

The following routines are found under the time domain.

ACF

Computes the autocorrelation function (ACF) for lags 0 through maximum lag, m , for a given series.

Input: The input box for ACF prompts the user for:

Series: ID number or full name of the series.

Max lag: Maximum lag m for ACF. Calculates from 0 to m .

Output: Displays the ACF, zero lag autocovariance, and the 95% significance.

Graph: Displays a graph of the ACF in the bottom window. The graph window has one additional menu option.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

PACF

Computes the partial autocorrelation function (PACF) for lags 0 to maximum lag, m , for a given series.

Input: The input box for PACF prompts the user for:

Series: ID number or full name of the series.

Max lag: Maximum lag m for PACF. Calculates from 0 to m .

Output: Displays the PACF and 95% significance.

Graph: Displays a graph of the PACF in the bottom window. The graph window has one additional menu option.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Variogram

Computes the variogram for lags 0 to m for a given series.

Input: The input box for variogram prompts the user for:

Series: ID number or full name of the series.

Max lag: Maximum lag m for variogram. Calculates from 0 to m.

Output: Displays the variogram.

Graph: Displays a graph of the variogram in the bottom window.

CCF

Computes the cross correlation function (CCF) for lags -m to m between two given series, series1(t) vs. series2(t-m). Positive lag indicates series1 correlation with the past of series2. while negative lag indicates series1 correlation with the future of series2. Plots the CCF as well as the ACF for each of the two series.

Input: The input box for CCF prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Max lag: Maximum lag m for CCF. Calculates from -m to m.

Output: Displays the results of the CCF and 95% significance. Does not display the ACF results.

Graph: Graphs of the CCF, ACF for the first series, and ACF for the second series are displayed in the bottom window. The graph window has two additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

All plots: Restores all three plots.

Autoregression

The autoregression routine evaluates autoregressions of increasing size from min. to max. order. Chooses as best, the autoregression with the smallest value of the model selection criterion chosen and displays a summary of its coefficients and forecasts, if any. Displays graphs of the predicted values, forecasts, residuals and the spectrum of the model.

Input: The input box for autoregression prompts the user for:

Series: ID number or full name of the series.

Min order: Minimum AR order to consider.

Max order: Maximum AR order to consider.

Forecasts: Number of forecasts. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include a constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Output: Displays a summary of the order, model selection criterion and first few parameters, of the models fitted. A summary of the best fitting model, with individual not global p-values for the parameters, appears next. If forecasts are requested, a summary of the forecast and forecast errors is presented. Lastly, the spectrum of the model.

Graph: Displays graphs of the predicted values, forecasts, residuals, and spectrum of the model in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has two additional menu options.

All Plots: Displays all three graphs.

Save: Saves residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

ARIMA

Evaluates a least squares version of an $ARIMA(p,d,q) \times (P,D,Q)_S$ model. ASTSA prints the current estimates for each iteration. If the estimates converge, a message box appears indicating that convergence has occurred. At this message, enter in a new convergence criteria or click Cancel to continue. If the number of allowed iterations are surpassed, another message box appears indicating that the number of allowed iterations has been surpassed. At this message, enter in any additional iterations, or click Cancel to continue. When the iterations are completed, a summary of the coefficients will be displayed. Four graphs will be displayed. These include predicted values, forecasts, residuals and the spectrum of the model.

Input: The input box for ARIMA prompts the user for:

Series: ID number or full name of the series.

p: AR order.

d: Amount of difference.

q: MA order.

P: Seasonal AR order.

D: Amount of seasonal difference.

Q: Seasonal MA order.

Season: Size of seasonal component.

Forecasts: Number of forecasts. 0 for no forecasts.

Use .1 guess (y/n): ARIMA uses either .1 as starting values or calculates starting values from ACF/PACF. Choose y to use .1 starting values. Choose n to use starting values based on ACF/PACF.

Selector (AIC, AICc, BIC, FPPEL, AICL): model selection criterion.

Output: Displays the iteration summary. A summary of the model, with individual not global p-values for the parameters. Model selection criterion for the model. If forecasts are requested, a summary of the forecast and forecast errors is presented. Lastly, the spectrum of the model.

Graph: Displays graphs of the predicted values, forecasts, residuals, and spectrum of the model in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has two additional menu options.

All Plots: Displays all three graphs.

Save: Save residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

ARIMA search

Searches for the best ARIMA model among a list of candidates. Then evaluates the best least squares version of an $ARIMA(p,d,q) \times (P,D,Q)_S$ model. ASTSA prints the current estimates for each iteration. When the best model is chosen and the iterations are completed, a summary of the coefficients will be displayed. Four graphs will be displayed. These include predicted values, forecasts, residuals and the spectrum of the model.

Input: The input box for ARIMA prompts the user for:

Series: ID number or full name of the series.

Min p: Minimum AR order to consider.

Max p: Maximum AR order to consider.

Min d: Minimum amount of difference to consider.

Max d: Maximum amount of difference to consider.

Min q: Minimum MA order to consider.

Max q: Maximum MA order to consider.

Min P: Minimum seasonal AR order to consider.

Max P: Maximum seasonal AR order to consider.

Min D: Minimum amount of seasonal difference to consider.

Max D: Maximum amount of seasonal difference to consider.

Min Q: Minimum seasonal MA order to consider.

Max Q: Maximum seasonal MA order to consider.

Min season: Minimum size of seasonal component to consider.

Max season: Maximum size of seasonal component to consider.

Forecasts: Number of forecasts. 0 for no forecasts.

Use .1 guess (y/n): ARIMA uses either .1 as starting values or calculates starting values from ACF/PACF. Choose y to use .1 starting values. Choose n to use starting values based on ACF/PACF.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Output: Displays the candidate models and the value of the model selection criterion. Then displays the iteration summary. A summary of the model, with individual not global p-values for the parameters. Model selection criterion for the model. If forecasts are requested, a summary of the forecast and forecast errors is presented. Lastly, the spectrum of the model.

Graph: Displays graphs of the predicted values, forecasts, residuals, and spectrum of the model in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has two additional menu options.

All Plots: Displays all three graphs.

Save: Save residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

Multiple regression

This univariate regression routine computes the multiple regression of one dependent series on n independent lagged series. Forecasting depends on the lags of the independent series. The number of forecasts can be no more than the smallest lag.

Input: The input box for multiple regression is of varying size. Initially prompts the user for:

Series: ID number or full name of the dependent series (Y).

No. of independent series: Number N of independent series (X).

Forecasts: Number of forecasts. Limited by the choice of lags for the independent series. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include a constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Once the number of independent series is chosen, the input box expands and prompts for:

Series: ID number or full name of the dependent series (Y).

No. of independent series: Number N of independent series (X).

Series 1: ID number or full name of the first independent series.

Lag: Time lag of first independent series. Must be positive

:

Series N : ID number or full name of the last independent series.

Lag: Time lag of last independent series. Must be positive.

Forecasts: Number of forecasts. Limited by the choice of lags for the independent series. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include a constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Output: Displays a summary of the model with individual not global p-values. Model selection criterion for the model. A summary of forecasts and forecast errors if requested. Multiple regression cannot compute more forecasts than the minimum lag of the independent series.

Graph: Displays graphs of the predicted values and residuals in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has two additional menu options.

All Plots: Displays all three graphs.

Save: Saves residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

State Space

State space modeling using the De Jong recursions. ASTSA prints the current estimates for each iteration. If the estimates converge, a message box appears indicating that convergence has occurred. At this message, enter in a new convergence criteria or click Cancel to continue. If the number of allowed iterations are surpassed, another message box appears indicating that the number of allowed iterations has been surpassed. At this message, enter in any additional iterations, or click Cancel to continue. When the iterations are completed, displays a summary of the coefficients, forecasts, and smoothed states. For the states, x-states menu option, displays one graph of the state and 95% error bounds. For the observed series y-obs menu option, displays three graphs predicted values, forecasts, and residuals.

The state space routine can handle structural models via structural matrices. These matrices should be composed of 0s and 1s. To hold a parameter fixed, use the 1 code. To update and estimate the parameter, use the 0 code.

Input: The input box for state prompts the user for:

Dimension of observation vector q (no. of series)

q :

Dimension of state vector p

p :

Forecasts, convergence, initial number of iterations

Forecasts: Number of forecasts. 0 for none.

Convergence: Relative convergence. Decimal less than .1.

Iterations: Initial number of iterations. Should be greater than 20.

M. V. code: Missing value code.

List of q observation vectors

Series 1: ID number of first series.

:

Series q : ID number of last series.

Initial x_0 values of dim p

X0 1: First component of x_0 .

:

X0 p : Last component of x_0 .

Prior covariance matrix for x_0 of dim $p \times p$

Row 1: First row of the covariance matrix. P columns of values separated by spaces.

:
 Row p: Last row of the covariance matrix. P columns of values separated by spaces.
 Transition matrix PHI of dim p x p
 Row 1: First row of the phi matrix. P columns of values separated by spaces.
 :
 Row p: Last row of the phi matrix. P columns of values separated by spaces.
 Structural matrix PHI of dim p x p (1=fix)
 Row 1: First row of structural phi matrix. P columns of 0s or 1s separated by spaces.
 :
 Row p: Last row of structural phi matrix. P columns of 0s or 1s separated by spaces.
 State covariance matrix Q of dim p x p
 Row 1: First row of the Q matrix. P columns of values separated by spaces.
 :
 Row p: Last row of the Q matrix. P columns of values separated by spaces.
 Structural matrix Q of dim p x p (1=fix)
 Row 1: First row of the structural Q matrix. P columns of 0s or 1s separated by spaces.
 :
 Row p: Last row of the structural Q matrix. P columns of 0s or 1s separated by spaces.
 Measurement matrix A of dim q x p
 Row 1: First row of the A matrix. P columns of values separated by spaces.
 :
 Row q: Last row of the A matrix. P columns of values separated by spaces.
 Structural matrix A of dim p x p (1=fix)
 Row 1: First row of the structural A matrix. P columns of 0s or 1s separated by spaces.
 :
 Row p: Last row of the structural A matrix. P columns of 0s or 1s separated by spaces.
 Observation covariance matrix R of dim q x q
 Row 1: First row of the R matrix. P columns of values separated by spaces.
 :
 Row q: Last row of the R matrix. P columns of values separated by spaces.
 Structural matrix R of dim p x p (1=fix)
 Row 1: First row of the structural R matrix. P columns of 0s or 1s separated by spaces.
 :
 Row p: Last row of the structural R matrix. P columns of 0s or 1s separated by spaces.

Output: Displays the iteration summary, the parameter estimates, the state vectors with errors terms, and if requested, forecasts with forecast errors.

Graph: Displays graphs of the predicted values, forecasts, residuals, for the first y-series, in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has three additional menu options.

X-states: List box of all the state vectors in the model. Double click on the list to view the graph. Click cancel to return to the current graph. Only one graph.

Y-observed: List box of all observation vectors in the model. Double click on the list to view the graph. Click cancel to return to the current graph. Three graphs

All Plots: Displays all three graphs.

Vector autoregression

Multivariate or vector version of autoregression. Denoted by VAR.

Input: The input box for VAR initially prompts the user for:

No. of series: Number of series q in the vector. Also dimension of the vector.

Min order: Minimum order to consider.

Max order: Maximum order to consider.

Forecasts: Number of forecasts. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include the constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Once the number of series is chosen, the input box expands and prompts for:

No. of series: Number of series q in the vector. Also dimension of the vector.

Series 1: ID number or full name of the first series.

:

Series q : ID number or full name of the last series.

Min order: Minimum order to consider.

Max order: Maximum order to consider.

Forecasts: Number of forecasts. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include the constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Output: Displays a summary of the models and Model selection criterion. A summary of the best fitting model by matrix and broken down into detail complete with individual p-values. Displays forecasts with errors if requested.

Graph: Displays graphs of the predicted values and residuals in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has three additional menu options.

Y-observed: List box of all observation vectors in the model. Double click to view the graph. Click cancel to return to the current graph.

All Plots: Displays all three graphs.

Save: Saves residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

Multivariate regression

The multiple regression routine evaluates the regression of Q dependent series on P independent lagged series. Forecasting depends on the lags of the independent series. The number of forecasts can be no more than the smallest lag.

Input: The input box for multivariate regression is of varying size. Initially prompts the user for:

No. of dependent series: Number q of dependent series (Y).

No. of independent series: Number p of independent series (X).

Forecasts: Number of forecasts. Limited by the choice of lags for the independent series. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include a constant. Choose n for no constant.

Selector (AIC, AICc, BIC, FPEL, AICL): model selection criterion.

Once the number of independent series is chosen, the input box expands and prompts for:

No. of dependent series: Number q of dependent series (Y).

Series 1: ID number or full name of the first dependent series.

:

Series q : ID number or full name of the last dependent series.

No. of independent series: Number p of independent series (X).

Series 1: ID number or full name of the first independent series.

Lag: Time lag of first independent series.

:

Series p : ID number or full name of the last independent series.

Lag: Time lag of last independent series.

Forecasts: Number of forecasts. Limited by the choice of lags for the independent series. 0 for no forecasts.

Include constant (y/n): Indicates whether a constant term should be included in the model. Choose y to include a constant. Choose n for no constant.

Output: Displays the results in matrix form as well as detailed form with p -values. Gives the model selection criterion for the model. If requested, gives forecasts with errors.

Graph: Displays graphs of the predicted values and residuals in the bottom window. For predictions, the thick line is the predicted values and the thin line is the actual data. The graph window has three additional menu options.

Y-observed: List box of all observation vectors in the model. Double click to view the graph. Click cancel to return to the current graph.

All Plots: Displays all three graphs.

Save: Saves residuals and/or the predicted values. A check mark appears next to the menu option when selected.

Residuals: Saves the residuals. Appends the residual series name to the end of the main ASTSA list.

Predicted: Saves the predicted values. Appends the predicted values series name to the end of the main ASTSA list.

Frequency Domain

Includes mostly descriptive routines. Also includes filtering. Most of the routines share four components, T', number of estimates, smoothing, and cosine taper. T' must be a power of two at least as great as the series length. If not, then ASTSA replaces T' with the next power of two greater than or equal to the series length. Number of estimates sets the number of spectral values displayed and graphed. The number of estimates must be a power of two less than T'. If not, then ASTSA uses 64 or half T', whichever is less. Due to limitations in Visual Basic®, T' cannot exceed 8192. If the series is longer than 8192, only the last 8192 observations are used. Smoothing sets the bandwidth or smoothing value L. L must be odd. If not, then ASTSA uses L+1. Cosine taper determines whether a cosine bell taper is applied to the data.

Frequency Domain	Transform
Spectrum	
Log spectrum	
Cepstrum	
Cross spectrum	
Coherence	
Impulse-response	
Filter1 - design a filter	
Filter2 - frequency response	

The following routines are found under frequency domain.

Spectrum

Computes the spectrum of the series using the fast Fourier transform.

Input: The input box for spectrum prompts the user for:

Series: ID number or full name of series.

T': Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the spectrum and the confidence band multipliers.

Graph: A graph of the spectrum will be displayed in the bottom window. The graph window has one additional menu option.

95% C.I.: Draws 95% confidence lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Log spectrum

Computes the log of the spectrum of the series using the fast Fourier transform.

Input: The input box for log spectrum prompts the user for:

Series: ID number or full name of series.

T': Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the log spectrum and the confidence band adders.

Graph: Displays a graph of the log spectrum in the bottom window. The graph window has one additional menu option.

95% C.I.: Draws 95% confidence lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Cepstrum

Computes the cepstrum of the series using the fast Fourier transform.

Input: The input box for cepstrum prompts the user for:

Series: ID number or full name of series.

T': Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the cepstrum.

Graph: Displays a graph of the spectrum in the bottom window.

Cross spectrum

Computes the gain, phase, quad spectrum, and cospectrum between two series.

Input: The input box for cross spectrum prompts the user for:

Series1: ID number or full name of first series.

Series2: ID number or full name of second series.

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the 't' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for 't'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the gain, phase, quad spectrum, and cospectrum.

Graph: A graph of the gain, phase, quad spectrum, and cospectrum will be displayed in the bottom window. The graph window has one additional menu option.

All plots: Displays all four plots.

Coherence

Computes the squared coherence between two series.

Input: The input box for coherence prompts the user for:

Series1: ID number or full name of first series.

Series2: ID number or full name of second series.

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the 't' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for 't'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the coherence only.

Graph: Displays two copies of the squared coherence and a graph of the spectrum for each series in the bottom window. The graph window has two additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

All plots: Displays all four plots.

Impulse response

Computes the impulse response or regression function for series 1 as input with series 2 as output and for the reverse, series 2 as input with series 1 as output. Positive indices indicate the regression coefficients relating the out series at time t with past values of the in series. Negative indices indicate the regression coefficients relating the out series with future values of the in series.

Input: The input box for impulse prompts the user for:

Series1: ID number or full name of first series.

Series2: ID number or full name of second series.

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Output: Displays the impulse response or regression function for series 1 as input and series 2 as output. Also displays the impulse response for series 2 as input and series 1 as output.

Graph: Two graphs of the impulse response functions will be displayed in the bottom window. One for each direction. The graph window has one additional menu option.

All plots: Displays both plots.

Filter1 - design a filter

Filter design allows the user to enter the frequency response of the filter and compute a filter which best matches the frequency response. A special band pass line in the input box allows quick creation of band pass filters. The format is l-u,v. l is the lower frequency, u is the upper frequency for the band, and v is the value of the frequency response across the band. Typing 0-.2,1 fills all the frequencies between 0 and .2 with 1.

Input: The input box for design a filter initially prompts the user for:

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Length of filter: Length of the desired filter. Adjusted to an even number.

[band input (l-u,v)]: Quick band input of desired frequency responses.

Input: Once the number of estimates is chosen, the input box expands and prompts for:

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Length of filter: Length of the desired filter. Adjusted to an even number.

[band input (l-u,v)]: Quick band input of desired frequency responses.

Freq .00000: Desired frequency response at frequency 0.

:

Freq .50000: Desired frequency response at frequency .5.

Output: Displays the filter coefficients and the actual frequency response for the filter.

Graph: Displays a plot of the filter coefficients, the desired frequency response, and the actual frequency response of the filter. The graph window has one additional menu option.

All plots: Displays all three plots.

Filter2 - frequency response of filter

Computes the frequency response of a filter.

Input: The input box for frequency response of filter initially prompts the user for:

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Length of filter: Length of the desired filter. Adjusted to an even number.

Input: Once the filter length is chosen, the input box expands and prompts for:

T: Underlying length of the series. Must be a power of 2 \geq the series length. Suggests a power of 2 based on the length of the longest series. Click on the t' to

view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T .

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Length of filter: Length L of the desired filter.

$a(-L/2)$: Value of the first filter coefficient.

:

$a(L/2)$: Value of the last filter coefficient.

Note that odd length filters have symmetric subscripts. Even length filters have an extra positive subscript.

Output: Displays the filter coefficients and the frequency response for the filter.

Graph: Displays a plot of the filter coefficients, and the frequency response of the filter. The graph window has one additional menu option.

All plots: Displays both plots.

3D

Dynamic 3D routines are primarily used to detect time varying processes. The series are broken into blocks by choosing the window size and overlap. The results of each block are displayed in a three dimensional graph. Scroll bars control the rotation and angle of the graph. Multiple 3D routines compare one series with all other series. All 3D graph windows share the following menus, Grid, Contours, Style, Surface, and Animate.

Grid: Sets the z scale axis min, max and increment as well as toggles the grid lines and portions of the cube surrounding the surface.

Z scale: Here the user can change the z scale minimum, maximum, and increment. An input box appears prompting the user to enter these values.

Z grid: Toggles the z grid lines.

XY grid: Toggles the xy grid lines on the bottom of the cube surrounding the surface.

Front block: Toggles drawing the front of the cube.

Back block: Toggles drawing the back of the cube.

Contours: Toggles the contour and density plot options. In a density plot, the four points on square of the grid of the surface plot are averaged and then shaded with smaller values darker than larger values.

Contours: Toggles the drawing of contour lines on the base of the cube under the surface.

Shaded contours: Similar to contours, but the bands are shaded with smaller values darker than higher values.

Shaded contours on surface: similar to shaded contours, but the drawing is on the surface itself.

Density plot: Toggles the density plot on the base of the cube under the surface.

Density plot on surface: Similar to density plot, but the drawing is on the surface itself.

Style: Determines the type of data displayed. Data can be averaged over the blocks or simply rescaled so that each window is rescaled to its own 0-1 scaling.

Raw data: Displays the original or raw data.

Average over s: Averages the data over one direction. Essentially a one two dimensional plot stretched over three dimensions.

Average over w: Averages the data over the other direction.

Remove average over s: Similar to average over s but the average is subtracted from the raw data.

Remove average over w: Similar to average over w but the average is subtracted from the raw data.

Rescale over s: Rescales each slice or block in one direction to its own 0-1 scale.

Rescale over w: Rescales each slice in the other direction to its own 0-1 scale.

Surface: Determines the surface type, lines or a grid. Four options. The current choice is check marked.

Spikes: Draws a spike over each point.

Lines over s: Draws parallel lines in one direction.

Lines over w: Draws parallel lines in the other direction.

Grid: Draws a grid over the surface. Lines in both directions.

None: No surface is displayed. This option makes viewing contours easier.

Solid surface: Toggles the hidden line removal. The sides of the surface appear transparent. Choose once to turn this feature on, choose again to turn this feature off. Gives the illusion that the surface is solid and not transparent by removing the parts of the graph that should not be visible. Sometimes the hidden line removal fails in which case the parts of the graph are missing. Choose this option again to turn off the hidden line routine. Try a different angle and try removing the lines again.

Solid block: Similar to solid surface but the sides of the surface appear solid as well.

Striped block: Similar to solid surface but the sides of the surface appear striped as well.

Height exaggeration: Sets the amount of height exaggeration, between 0 and 1, in the surface. The higher the value the taller peaks appear. An input box appears where the user enters the amount of height exaggeration scaling as a number between 0 and 1.

Animate: Causes the surface to rotate by d increments in degrees. An input box appears where the user can specify the rate of rotation in degrees. Click on animate again to stop.

3D Transform
Dynamic ACF
Dynamic PACF
Dynamic CCF
Dynamic Spectrum
Dynamic log spectrum
Dynamic Cepstrum
Dynamic Coherence
Multi CCF
Multi Coherence

Routines found under 3D

Dynamic ACF

Breaks the series into blocks and computes the autocorrelation function (ACF) for lags 0 to m for each block. The blocks may overlap. The resulting ACFs are stacked to form a surface plot. Useful for studying whether the structure of the series changes over time.

Input: The input box for dynamic ACF prompts the user for:

Series: ID number or full name of the series.

Lag: Maximum lag m for each ACF. Calculates from 0 to m .

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the ACF calculations.

Output: Displays the results of all the ACFs including 95% significance in the top window.

Graph: A 3D surface plot of all the ACFs is displayed in the bottom window. A spike surface is initially displayed. Line and grid surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Dynamic PACF

Breaks the series into blocks and computes the partial autocorrelation function (PACF) for lags 0 to m for each block. The blocks may overlap. The resulting PACFs are stacked to form a surface plot. Useful for studying whether the structure of the series changes over time.

Input: The input box for dynamic PACF prompts the user for:

Series: ID number or full name of the series.

Lag: Maximum lag m for each PACF. Calculates from 0 to m .

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the PACF calculations.

Output: Displays the results of all the PACFs including 95% significance in the top window.

Graph: A 3D surface plot of all the PACFs is displayed in the bottom window. A spike surface is initially displayed. Line and grid surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Dynamic CCF

Breaks the two series into blocks and computes the cross correlation function (CCF) for lags $-m$ to m for each block. The blocks may overlap. The resulting CCFs are stacked to form a surface plot. Useful for studying whether the relationship between two series changes over time.

Input: The input box for dynamic CCF prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Lag: Maximum lag m for each CCF. Calculates from $-m$ to m .

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the CCF calculations.

Output: Displays the results of all the CCFs including 95% significance in the top window.

Graph: A 3D surface plot of all the CCFs is displayed in the bottom window. A spike surface is initially displayed. Line and grid surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Dynamic spectrum

Breaks the series into blocks and computes the spectrum for each block. The blocks may overlap. The resulting spectra are stacked to form a surface plot. Useful for studying whether the series changes over time.

Input: The input box for dynamic spectrum prompts the user for:

Series: ID number or full name of series.

T: Underlying length of the series. Must be a power of 2 \geq the window length.

Suggests a power of 2 based on the length of the longest series. Click on the 't' to

view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T' .

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the spectrum calculations.

Output: Displays the results of all the spectra.

Graph: A 3D surface plot of all the spectra is displayed in the bottom window. A grid surface is initially displayed. Line surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Dynamic log spectrum

Breaks the series into blocks and computes the log of the spectrum for each block. The blocks may overlap. The resulting log spectra are stacked to form a surface plot. Useful for studying whether the series changes over time.

Input: The input box for dynamic log spectrum prompts the user for:

Series: ID number or full name of series.

T' : Underlying length of the series. Must be a power of 2 \geq the window length. Suggests a power of 2 based on the length of the longest series. Click on the t' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T' .

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the spectrum calculations.

Output: Displays the results of all the log spectra.

Graph: A 3D surface plot of all the log spectra is displayed in the bottom window. A grid surface is initially displayed. Line surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Dynamic cepstrum

Breaks the series into blocks and computes the cepstrum for each block. The blocks may overlap. The resulting cepstra are stacked to form a surface plot. Useful for studying whether any delay structures in the series change over time.

Input: The input box for dynamic cepstrum prompts the user for:

Series: ID number or full name of series.

T: Underlying length of the series. Must be a power of 2 \geq the window length. Suggests a power of 2 based on the length of the longest series. Click on the 't' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for 't'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the cepstrum calculations.

Output: Displays the results of all the cepstra.

Graph: A 3D surface plot of all the cepstra is displayed in the bottom window. A grid surface is initially displayed. Line surfaces are available.

Dynamic Coherence

Breaks the two series into blocks and computes the squared coherence for each block. The blocks may overlap. The resulting coherence are stacked to form a surface plot. Useful for studying whether the relationship between two series changes over time.

Input: The input box for dynamic coherence prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

T: Underlying length of the series. Must be a power of 2 \geq the longest series length. Suggests a power of 2 based on the length of the longest series. Click on the 't' to

view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for t' .

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T .

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

Window length: Size for each of the blocks.

Overlap: Sets the overlap for the blocks. Choose 0 for no overlap. Choose a negative number to overlap the blocks. Choose a positive number to leave gaps between the blocks. The data in the gaps is not used for the coherence calculations.

Output: Displays the results of all the squared coherence.

Graph: A 3D surface plot of all the coherence is displayed in the bottom window. A grid surface is initially displayed. Line surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

MultiCCF

Computes the cross correlation function (CCF) for lags $-m$ to m between a given series, as $series1(t)$, and all other series, as $series2(t-m)$ including itself. Positive lag indicates $series1$ correlation with the past of $series2$. while negative lag indicates $series1$ correlation with the future of $series2$. The resulting CCFs are stacked to form a surface plot.

Input: The input box for multiCCF prompts the user for:

Series: ID number or full name of the series.

Max lag: Maximum lag m for each CCF. Calculates from $-m$ to m .

Output: Displays the results of all the CCFs including 95% significance in the top window.

Graph: A 3D surface plot of all the CCFs is displayed in the bottom window. A spike surface is initially displayed. Line and grid surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

MultiCoherence

Computes the squared coherence between the selected series and all other series. The resulting coherence are stacked to form a surface plot.

Input: The input box for multicoherence prompts the user for:

Series: ID number or full name of first series.

T': Underlying length of the series. Must be a power of 2 \geq the longest series length. Suggests a power of 2 based on the length of the longest series. Click on the 't' to view a list of powers of 2. Click on a power of 2 in the list to choose this power of 2 for 't'.

Number of estimates: Number of spectral estimates used to create the graph. Also must be a power of 2 and less than T'.

Smoothing: Amount of smoothing (averaging) used to generate each displayed spectral estimate.

Cosine taper (y/n): Toggles the cosine tapering option. Choose y to taper the data. Choose n to use the raw data.

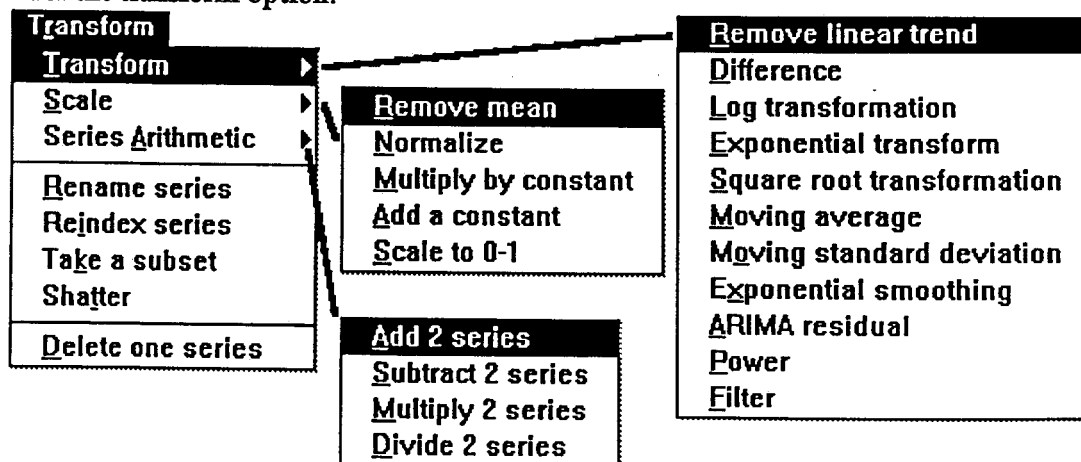
Output: Displays the results of all the squared coherence.

Graph: A 3D surface plot of all the coherence is displayed in the bottom window. A grid surface is initially displayed. Line surfaces are available. In the spike view, toggling the 95% significance lines causes the non significant values to vanish. The graph window has three additional menu options.

95% sig: Draws 95% significance lines on the graph. Click once to draw the lines, click again to erase the lines. Note that this is individual not global significance.

Transform

Includes routines for transforming, scaling the data. The original data remains unchanged. The transformed data is appended to the end of the main ASTSA list. The new series uses the original series name plus the transformation type. The nine transformations found under the transform sub menu. The five scaling options found under the scale sub menu. The following routines are found under the transform option:



Under transform, we have the following options.

Remove linear trend

Removes linear trend from a given series. A message box will display the slope and y-intercept of the linear trend. Denoted by the "linear".

Input: The input box for remove linear prompts the user for:
Series: ID number or full name of the series.

Difference

Applies a d lag difference to the series. Denoted by "D#". Where # is the d lag difference used.

Input: The input box for difference prompts the user for:
Series: ID number or full name of the series.
D: lag for the difference.

Log transform

Applies a log transformation to the series. If the series contains 0 or negative values, the minimum-1 will be subtracted from each point and then the log is taken. Denoted by "log". If the series contained non-positive values, denoted by "-min+1 log".

Input: The input box for log transform prompts the user for:
Series: ID number or full name of the series.

Exponential transformation

Transform the data by creating $x_t = \exp(y_t)$, where y_t is the original series. Denoted by "exp"

Input: The input box for exponential transformation prompts the user for:
Series: ID number or full name of the series.

Square root transform

Applies a square root transformation to the series. If the series contains negative values, the minimum will be subtracted from each point and then square root is taken. Denoted by "sqrt". If the series contains negative values, denoted by "-min sqrt".

Input: The input box for log transform prompts the user for:
Series: ID number or full name of the series.

Moving average

Applies a centered moving average to the series. Denoted by "MA#". Where # is the size of the moving average.

Input: The input box for moving average prompts the user for:
Series: ID number or full name of the series.
Order: Size of the moving average.

Moving standard deviation

Applies a centered moving standard deviation to the series. Denoted by "MSD#". Where # is the size of the moving standard deviation.

Input: The input box for moving standard deviation prompts the user for:

Series: ID number or full name of the series.

Order: Size of the moving standard deviation.

Exponential smoothing

Exponentially smoothes the series. Denoted by "ES#". Where # is the smoothing term.

Input: The input box for exponential smoothing prompts the user for:

Series: ID number or full name of the series.

Factor: Smoothing term between 0 and 1.

ARIMA residual

Returns to residuals for a specified ARIMA model. All parameter values must be supplied by the user. Denoted by "ARIMA". The input box for ARIMA residuals is of varying size and expands depending on the order of the model.

Input: The input box for ARIMA residuals prompts the user for:

Series: ID number or full name of the series.

p: Order p of the autoregression.

d: Order d of the difference.

q: Order q of the moving average.

P: Order P of the seasonal autoregression.

D: Order D of the seasonal difference.

Q: Order Q of the seasonal moving average..

Season: Length of the seasonal component.

phi(1): AR parameter value. Appears if $p > 0$.

:

phi(p): AR parameter value. Appears if $p > 0$.

theta(1): MA parameter value. Appears if $p > 0$.

:

theta(q): MA parameter value. Appears if $p > 0$.

PHI(1): Seasonal AR parameter value. Appears if $p > 0$.

:

PHI(P): Seasonal AR parameter value. Appears if $p > 0$.

THETA(1): Seasonal MA parameter value. Appears if $p > 0$.

:

THETA(Q): Seasonal MA parameter value. Appears if $p > 0$.

Power

Raises all values in the series to the specified power. If an error occurs due to negative values in the series. The minimum will be subtracted and then the series raised to the power. Denoted by "power#". Where # is the power term.

Input: The input box for power prompts the user for:
Series: ID number or full name of the series.
Power: Power to raise each value.

Filter

Applies a filter to the series. One of the filters routines must be run prior to this routine. Denoted by "filtered".

Input: The input box for exponential smoothing prompts the user for:
Series: ID number or full name of the series.

Under scale, we have the following options.

Remove mean

Subtracts the mean from a given series. A message box will display the mean of the series. Denoted by "mean".

Input: The input box for remove mean prompts the user for:
Series: ID number or full name of the series.

Normalize

Normalizes the series by subtracting the series mean and dividing by the series standard deviation. Denoted by "norm".

Input: The input box for normalize prompts the user for:
Series: ID number or full name of the series.

Multiply by constant

Multiplies all values in the series by a specified constant. Denoted by "x #". Where # is the constant.

Input: The input box for multiply prompts the user for:
Series: ID number or full name of the series.
Constant: Value of constant to multiply with series.

Add a constant

Adds all values in the series with a specified constant. Denoted by "+ #". Where number is the constant.

Input: The input box for add prompts the user for:
Series: ID number or full name of the series.
Constant: Value of the constant to add to the series.

Scale to 0-1

Scales the series to 0-1 by subtracting the minimum and dividing by the range. Denoted by "scaled".

Input: The input box for scale prompts the user for:

Series: ID number or full name of the series.

Under series arithmetic, we have the following.

Add 2 series

Adds the values of two series together. Denoted by name of series 1 + name of series 2.

Input: The input box for add prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Subtract 2 series

Subtracts the values of the second series from the values of the first series. Denoted by name of series 1 - name of series 2.

Input: The input box for subtract prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Multiply 2 series

Multiplies the values of two series together. Denoted by name of series 1 x name of series 2.

Input: The input box for multiply prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Divide 2 series

Divides the values of the first series by the values of the second series. Denoted by name of series 1 / name of series 2.

Input: The input box for divide prompts the user for:

Series1: ID number or full name of the first series.

Series2: ID number or full name of the second series.

Rename series

Allows the user to rename a series.

Input: The input box for rename prompts the user for:

Series: ID number or full name of the series.

New name: New name for the series.

Reindex series

Reindexes the first active point as index 1. Many routines, particularly the subset routine, delete the end points off the series. These end points become inactive. The two series diagnostics and models require the series to share common active indices. Reindexing allows the series to begin at index 1. Denoted by "indexed".

Input: The input box for reindex prompts the user for:

Series: ID number or full name of the series.

Take a subset

Causes ASTSA to use only a subset of the series in future calculations. Denoted by "sub#-#". Where the first # is the starting time index. The second # is the stopping time index.

Input: The input box for subset prompts the user for:

Series: ID number or full name of the series.

Start: Starting time index.

Stop: Stopping time index.

Shatter

Breaks the series into multiple subsets of length equal to the window.

Input: The input box for shatter prompts the user for:

Series: ID number or full name of the series.

Window: Length of the subsets.

Overlap: Amount of overlapping from one subset to the next.

Delete

Deletes a series from ASTSA. The series must first be selected from the main list by single clicking on the desired series. This selected series will be highlighted. An "are you sure?" message box appears. Choose yes to delete the series. Choose no to cancel and return to the ASTSA main menu.

Uninstall



UnInstall

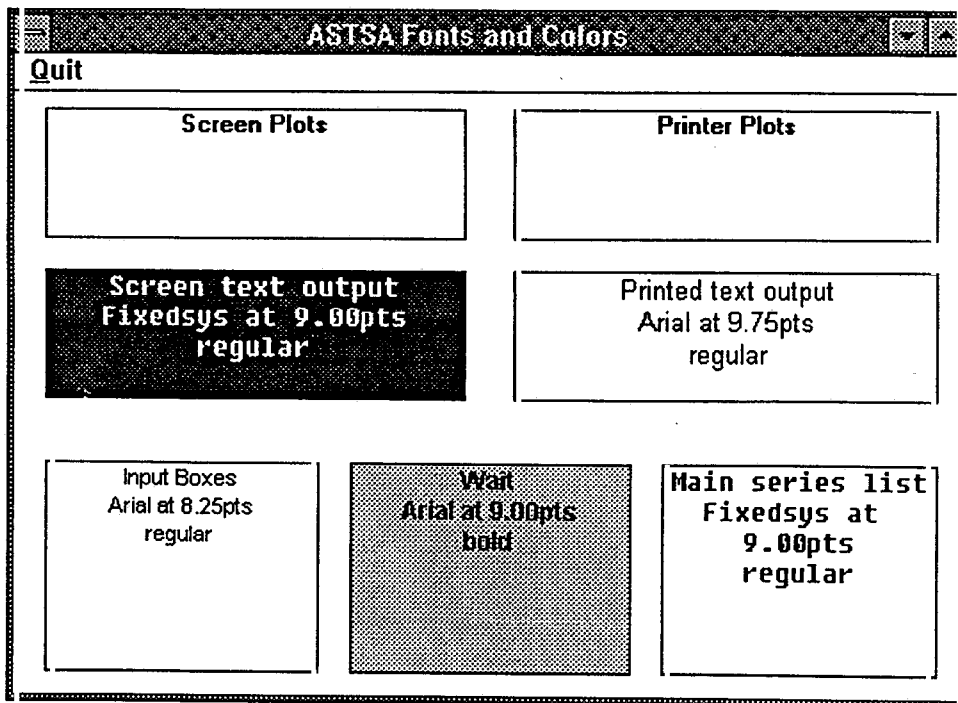
ASTSA may be uninstalled (removed) from the computer at any time. Double click on the uninstall icon in the WinASTSA group. An "Are you sure?" prompt appears. Choose delete to uninstall ASTSA. Choose cancel to cancel. The uninstall routine removes all ASTSA files from disk, removes the ASTSA sub directory, and deletes the WinASTSA group.

Fonts and Colors



Fonts, font sizes, and colors used by ASTSA may be changed. Changing the font size is useful if Windows is running at high resolution - beyond 800 x 600 super VGA. At higher resolutions, the words shrink until they are difficult to read. Increasing the font size can help increase readability. If fonts are added or removed from Windows, ASTSA may fail to run. Use the fonts program to select new fonts that are installed under Windows. Proportionally spaced fonts are of higher quality but distort the columns of tables. Non proportionally spaced fonts, such as fixedsys, display tables well but are of lower print quality. Non proportionally spaced fonts should be used for on screen display of tables, such as text output found below. For users with monochrome or LCD laptop displays, the color choices may be changed. The Fonts & Colors program changes these choices. ASTSA supports color printers. Fonts & Colors allows the user to change printing colors and styles to take advantage of color printers. Each of the seven parts displays the current font choice, using the current font, the part name, current font, font size and bold status printed in the current colors. Choose one of the command boxes to change the fonts for this part of ASTSA. Click on quit to quit the Fonts & Colors program.

The main Fonts & Colors window with the seven parts of ASTSA.

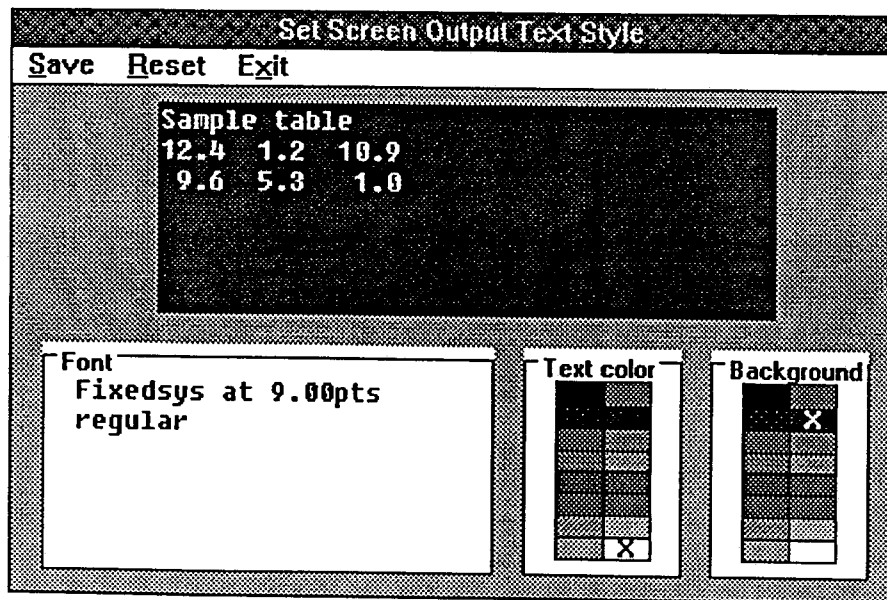


Arial and fixedsys are standard Windows 3.1 fonts. Arial is a high quality true type font. Fixedsys is a lower quality font good for displaying tables. We describe the graph output, text output, and print graphs in more detail. Clicking on the boxes calls the detail windows where the changes are made.

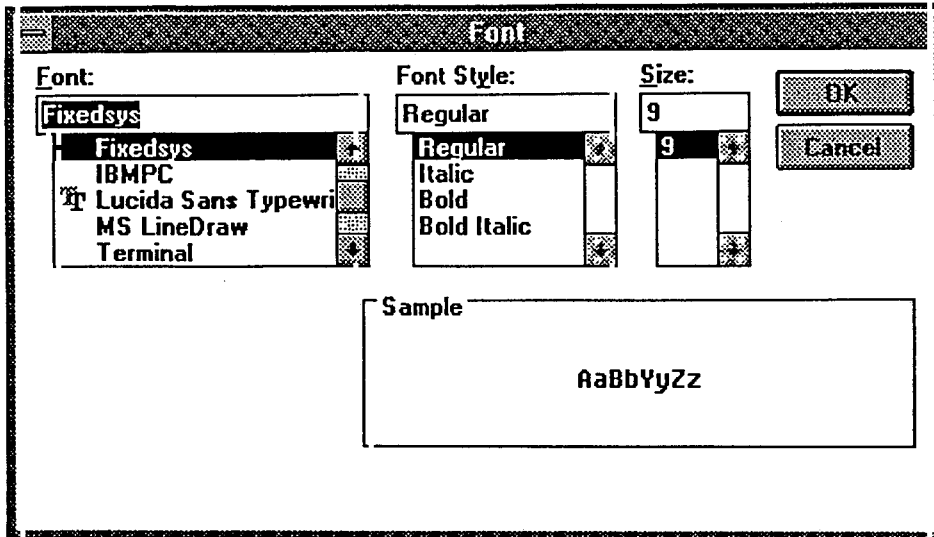
For all seven parts, a window appears with a font list, a font size list, bold option. Some of these windows have extra options and are described below in more detail. All seven share three menu items save, reset, and exit. Save updates the current choices and closes the window returning the user to the main Fonts & Colors window. Reset restores the original settings. Exit closes the window without saving the current choices.

Four of the parts display text on the screen. Main series list contains the font and color choices for the list of data series in the main ASTSA menu window. Screen text output contains the font and color choices for the output window. Wait box contains the settings for the status windows. These windows appear when ASTSA is working on a lengthy procedure. Input windows contains font and color choices for the input boxes. Note that the left side text is always bold faced. Print text output has similar choices to the other five text parts but has no color options. Print text output determines how text is printed on the printer.

Text output allows the user to change the font and colors of text displayed in a window. The output text styles windows has six parts plus the two menu items. The top text box displays the current font and color choices and displays a 2 x 3 table. The decimals in the table should line up. Click on the lower left box to view the font dialog box, shown below. The text colors displays the sixteen color choices for the text. The background displays the sixteen choices for the background. These colors cannot be the same.



An example of the font and color choice window described above.



The font dialog box. Choose font, size, and style.