Chapter 3 Window Types

Chapter Overview

This chapter describes the operations for each of the different types of windows that appear in NCSA DataScope—text windows, image windows, and notebook windows. These windows are used for displaying numbers as text, viewing images of different types, entering notes and formulas, and performing calculations to derive new datasets.

The Text Window

When you load a dataset, NCSA DataScope creates a text window with row and column labels, scroll bars, and a portion of the array printed as floating-point numbers.

You can control the size of this window using the size box at the lower-right corner of the window, scroll through the array using the horizontal and vertical scroll bars, and close the window by clicking in the close box at its upper-left corner. Similar in nature to spreadsheet programs, the text window allows you to select and copy to the Clipboard the array of numbers or portions of it. For more information about selecting numbers from the text window, see this chapter's section entitled "Selecting Regions."

Specifying Text Window Characteristics

You can control several of the text window characteristics using the Attributes dialog box. To access this dialog box, shown in Figure 3.1, choose Attributes from the Numbers menu or press ⌘-A.

Figure 3.1 Attributes Dialog Box

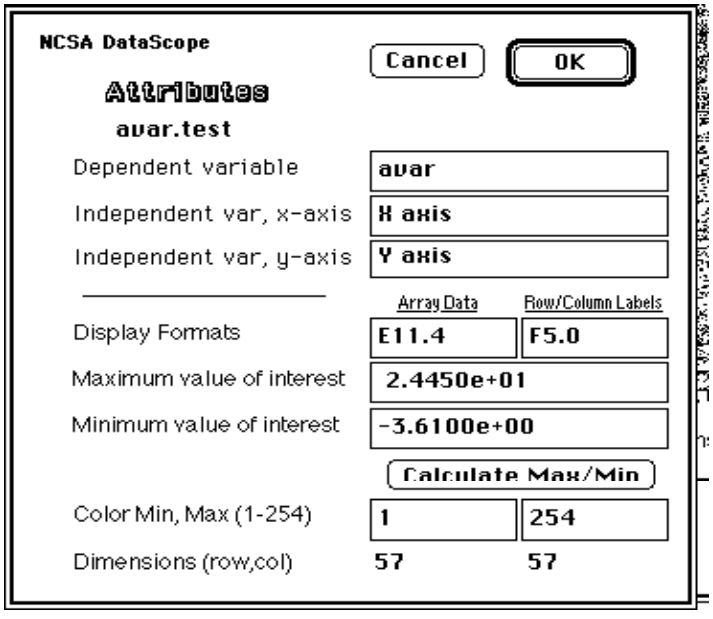

Using the Attributes dialog box, you can change the variable names for the active dataset, control the printing display format for numbers in the text window, specify alternate maximum and minimum values of interest, calculate the actual maximum and minimum values, and determine what range of colors from the 254-color palette will be used to represent the dataset in your image windows.

After making settings in the Attributes dialog box, click OK to enable your changes or click Cancel to undo them.

The fields in the Attributes dialog box have the following characteristics:

- The filename displayed under the Attributes heading (avar.test in Figure 3.1) is the original name of the dataset file. You cannot change this field from within the Attributes dialog box; it changes automatically when you rename and save the file.

- The three boxes labeled Dependent variable; Independent var, x-axis; and Independent var, y-axis are associated with the dependent variable and the respective axis variable. The dependent variable name is used for the text window title and is also used in calculations, as described in Chapter 4, "Notebook Calculations." To change the names of the variables, enter the variable names in the appropriate boxes.

- The two boxes labeled Display Formats control the printing display format for numbers in the text window. The default printing format is E11.4 for the data values and F5.0 for the column and row labels. To specify a different E- or F-format for the array or column and row labels, enter the desired format in the the box labeled Array Data or Row/Column Labels, respectively.

- The boxes labeled Maximum value of interest and Minimum value of interest control which values are considered the minimum and maximum values for image generation. When you load a dataset, the maximum and minimum values are calculated for you and placed here.

To ignore *outliers*, numbers which are larger or smaller than the important portion of the data, during image generation enter the maximum and minimum values of the pertinent data range in the text boxes labeled "Maximum and Minimum value(s) of interest" (Figure 3.1). Outliers are not actually removed from the dataset; they are just ignored during image generation.

To calculate the actual maximum and minimum values of the dataset, click the button labeled Calculate Max/Min. This overrides the maximum and minimum values of interest currently in the fields of the Attributes dialog box, displaying the actual largest and smallest values in the dataset in these fields

instead. When computing raster displays (polar, generated, interpolated), the data will be bounded to be within this range.

- The boxes labeled Color Min, Max (1-254) control the color range selection, which affects image generation only. From a maximum possible palette of 254 colors, only colors from this range are used by the image generation process. The default values are 1 and 254.

- The fields labeled Dimensions (row, col) display the actual row and column dimensions of the dataset, respectively, and cannot be changed using the Attributes dialog box.

Selecting Regions

All of the display windows in NCSA DataScope support the selection of a region of data values.

To select a single data value, click on the value in the text window or on a point in an image window.

To select a region of values in a text or image window, drag from one corner to the opposite corner of the desired region and release the button.

A key feature of NCSA DataScope is the synchronization of selection regions between windows. When you select a region in the text window, the corresponding region in each associated image window is outlined with a rectangle if this menu item is checked. (Associated images are those which were generated from the dataset.) Conversely, when you select a region in the image window by clicking or dragging out a rectangle, the text window highlights the selection region for the corresponding data values, scrolling to display the selection region in the center of the window.

Extracting Datasets

In Chapter 1 you learned that you can generate new datasets using the notebook window calculations, which are more fully discussed in Chapter 4, "Notebook Calculations." You can also create a new dataset by extracting it from an existing dataset:

1. Select the desired region of values in a text or image window associated with the source dataset.

   NOTE: Your selection region must contain at least two rows and two columns.

2. Choose Extract Selection from the Numbers menu. This pulls the data from the selection region of the current dataset and creates a new dataset with those values. A new text window displaying the extracted data appears on the screen.

NOTE: Unless you select the entire window as the selection region, this new subset has smaller row and column dimensions

than the original. The necessary values for the row and column scales are extracted from the original row and column scales.

The Image Windows

With an array in memory, NCSA DataScope can calculate and generate images from the dataset and create windows in which to display them. NCSA DataScope generates three types of images:

- simple, scaled color image
- interpolated (smoothed) image
- polar image

The supported types of image generation are described in the following sections.

Simple Color Image

The simplest and quickest imaging operation is that which produces the basic, scaled color image.

To generate a simple, scaled image for any dataset that is already loaded:

1. Activate the dataset by clicking on one of the windows associated with it.

2. Select Generate Image from the Image menu or press ⌘-G.

The algorithm used to generate the image works as follows. Calculating the difference between the maximum and minimum values of interest yields the range of values of interest. Scaling each value of the dependent variable with the following formula while the color minimum and maximum values are the default 1 and 254, for example, produces a linearly scaled set of numbers between 1 and 254. For example,

colorvalue = *1 + 254**[(value – minimum)/range]

Using the Macintosh color palette manager, NCSA DataScope assigns the values from the minimum color value (in this case, 1) to the range of hardware color values by default (in this case, 254). This produces a grey-scale or color raster image of the scaled data. You can change this range of colors in the Attributes dialog box (see "Specifying Text Window Characteristics").

Display of the data is achieved by a straightforward duplication of pixels. In the one-to-one case —where the number of pixels used to display the image matches the number of data points in the array—each pixel is set to the color for its corresponding data point.

When the data array has fewer points than the number of pixels to display, each data point is blown up to a rectangle of pixels. For

example, one data value may become a square of 3 x 3 pixels which are all set to the same color. This creates a chunky or jagged effect. To get smoother enlargements, use the interpolation method of image generation, which is discussed in the following section.

Interpolated (Smoothed) Image

The default color scaling for an interpolated image is on a linear 1 to 254 scale as for a scaled color image. ("Specifying Text Window Characteristics" discusses how you can change this scale.) For one-to-one display, when the number of pixels matches the number of data points, the images produced by the two methods are identical. When the number of pixels in the image outnumbers the number of data values in the data array (for example, when you increase the size of the image using the Interpolate Size command), each pixel of the image must be interpolated between multiple data values before it is scaled on the linear color scale. For information regarding changing the size of an image, see "Controlling the Image Size."

NOTE: Only one interpolated, generated or polar image, will be displayed for any one dataset at any one time.

To generate an interpolated image for any dataset that is already loaded:

1. Activate the dataset by clicking on one of the windows associated with it.

2. Select Interpolated Image from the Image menu or press ⌘-I.

NOTE: To create an image for a *unevenly spaced grid*, a data array without regularly spaced row and column scales, you should use the interpolation method. This method takes into account the X and Y distances between data values when performing the interpolation. For more detailed information about unevenly spaced grids, see the section under "Two-Dimensional Datasets" in Chapter 2.

Polar Image

The default color scaling for polar image generation is also on a linear 1 to 254 scale; however, the pixels are placed into the image according to a polar interpretation of the row and column scaling information. The row scale is treated as a set of radius values, and the column scale as a set of angle values in the range 0 to 2π radians. By default, the angle is measured clockwise from the horizontal axis, with the origin at the center of the window.

To generate a polar image for any dataset that is already loaded:

1. Activate the dataset by clicking on one of the windows associated with it.

2. Select Polar Image from the Image menu or press ⌘-P.

Note that when you select a region of data values from a text window, the region is realistically highlighted in the associated polar image with angles and radii, rather than a cartesian rectangle.

Image Size

To specify a size for a basic scaled or interpolated image to be generated:

1. Choose Image Size or Interpolate Size, respectively, from the Image menu. The Image Size Selection dialog box appears (see Figure 3.2).

2. Specify the width or height of the image by clicking the up or down arrow Width Expansion Factor or Height Expansion Factor, respectively, or by entering a value into the text box between the arrows.

   or

   Enter positive, nonfractional values in the boxes that appear at the right and bottom edges of the picture to specify the actual pixel height and width of the image.

   NOTE: NCSA DataScope only allows integer expansion factors.

3. Click OK or press RETURN to set the specifications for the next scaled or interpolated image you generate.

An image you generate using the Generate Image or Interpolate Image commands is calculated according to the specifications you last made in the Image Size Selection dialog box.

Figure 3.2 **Image Size Selection Dialog Box**

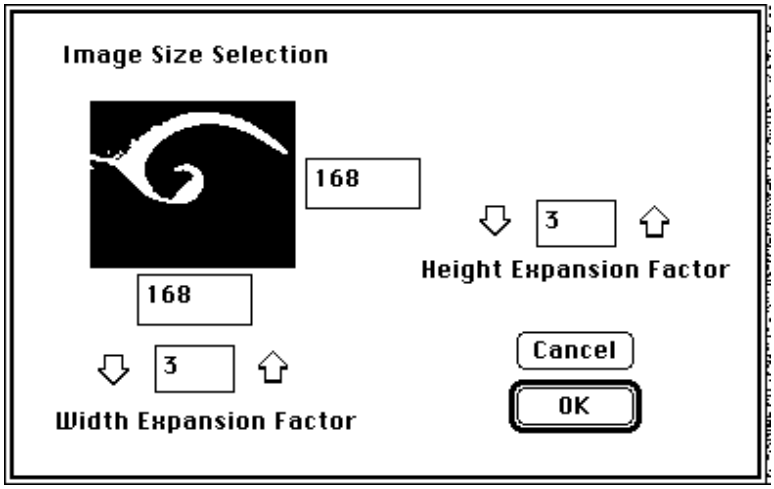

Selecting the size of a polar image involves a few more options. You can choose the direction of the zero angle, the portion of the circle to view, and an expansion factor that helps determine the final size of the viewing window. Polar data is defined on a circle with angles (columns) from zero to 2π, and radius values (rows) from zero to the maximum radius.

To specify the size of a polar image:

1. Choose Polar Size from the Image menu. The Polar Size Selection dialog box appears (see Figure 3.3).

2. Choose a direction for the zero angle by clicking one of the four radio buttons labeled North, East, South, and West, respectively. Angles proceed from this point in clockwise direction.

3. Drag out a rectangle within the picture frame to select a viewing region. When you generate the image, the relative portion of the image inside your viewing rectangle is visible in the window. The size of the circle in this window represents the scaled size of the maximum radius value (row label) given in the dataset.

4. Specify an expansion factor to be applied to your viewing rectangle, by clicking the up or down arrow labeled Expansion Factor or by entering a positive, nonfractional value in the text box between these arrows. For your reference, the numbers labeled Window Width and Window Height indicate the size of the resulting window.

   NOTE: NCSA DataScope only allows integer values.

5. Click OK or press RETURN to set the specifications for the next image generated.

Figure 3.3 **Polar Size Selection Dialog Box**

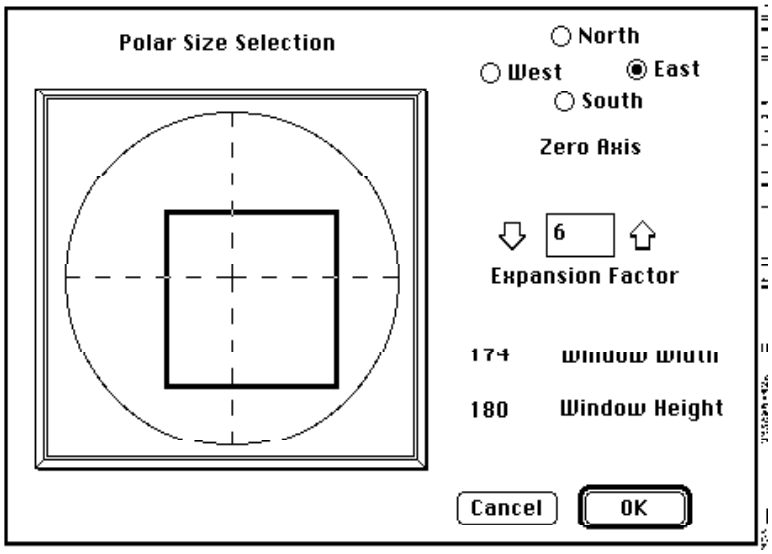

Color Palettes

The term *color palette* refers to the colors that are assigned to the range of values which NCSA DataScope creates for the pixels. This range is 1 to 254 by default, but you may change it using the Attributes dialog box (see "Specifying Text Window Characteristics"). The hardware supports 256 colors from a total selection of over 16 million colors. The color palette is the assignment of the scaled data values (1 to 254 by default) to the colors that you see on the screen.

NOTE: The values 1 and 256 are used by the system.

NCSA DataScope displays images with either a grey-scale palette or the most recently used palette by default. In the default grey-scale palette, the level of grey, or intensity of white, in the image shows the relative magnitude of the data value. This is the simplest mapping of colors (levels of grey in this case) to the data values, but many datasets can benefit from alternative mappings—a spectrum from blue to red, for example.

To assign a new set of colors to the current image window:

1. Select Load Palette from the File menu or press ⌘-L. A directory dialog box appears.

2. Select and open any HDF file which contains a palette. NCSA DataScope always loads the first palette in the file.

Several sample palettes are provided with NCSA DataScope, including a spectrum sample. Refer to the *NCSA HDF Calling Interfaces and Utilities* documentation for information regarding creating your own palettes; or use NCSA PalEdit, a palette creation program for the Macintosh.

Image Generation

A 400 x 400 interpolated image might take 30 seconds to complete, whereas the non-interpolated image only takes a few seconds. The standard image generation is dependent only on the size of the original data; the interpolated image, on the other hand, is more dependent on the size of image being generated, for every displayed pixel is separately calculated. Polar image generation primarily depends on the source data size, but the target image size is also a factor.

In all cases, first an empty window and then a watch cursor appears while the image is being created. The window is periodically updated as the image is generated, displaying the portion of the image that has been calculated so far.

Stopping Image Generation

All three methods of image generation update the screen with partial results as they are calculated. To stop the image generation before it is finished, press ⌘ and type a period (.). When you enter ⌘-PERIOD, the image generation process stops and the window remains in its current state with a partially generated image.

NOTE: Leaving a window in this state does not affect the operation of other actions; it is possible to work with only a portion of an image if you do not want to wait for the whole image to generate.

Window Synchronization

One of the most powerful features of NCSA DataScope is its ability to synchronize multiple datasets across windows. These might be two, three, or four variables from the same simulation, or multiple time-steps from a simulation run. The only requirement is that the number of rows and the number of columns of the datasets are identical.

To invoke the synchronization of selection regions, select Synchronize from the Numbers menu. A checkmark appears by the command to indicate that the Synchronize option is activated.

When the Synchronize option is active, a selection in one window produces matching selections in all text window datasets that have the same row and column dimensions and all image windows that were generated from those datasets. When you are done, select the Synchronize command again to turn the feature off and remove the checkmark.

NOTE: While the Synchronize option is on, any selection activity is reflected in all synchronized windows.

Image windows are always synchronized with their originating text windows regardless of the setting of the Synchronize option. The Synchronize command enhances this feature by synchronizing multiple, separate datasets.

The Notebook Window

In NCSA DataScope, every dataset has an associated *notebook*. Initially, this notebook is empty. You may enter any text in the NCSA DataScope's notebook window and save that information with the dataset. If the dataset was previously saved with a notebook, then that notebook window appears automatically when you reopen the HDF file.

If there is no notebook window or if it has been removed from the screen, select See Notebook from the Numbers menu or press ⌘-N.

The notebook for the current dataset appears on the screen. If it is already on the screen, it is brought to the front.

The notebook window is a standard Macintosh text window, so you can edit text in the notebook window as you would in a Macintosh word processor. You can also copy and paste between notebook windows and other text programs under MultiFinder. When you save the dataset, the entire contents of the notebook window are saved with it.

As described in Chapter 4, "Notebook Calculations," the notebook window can be used to enter formulas and perform calculations. You can save formulas in the notebook window, along with your comments and observations, for later use.