4 Visualizing Datasets

SciAn visualizes datasets using **visualization objects**, which visualize one or more datasets. There are several ways to make a visualization object within SciAn. The easiest is to select the datasets to visualize in the Datasets window and press the Visualize button or choose Visualize from the Object menu. SciAn will examine the data and choose a visualization technique and default settings.

Visualization windows and spaces

Visualization objects must be displayed in visualization windows to be seen. A visualization window can hold any number of visualization objects. Figure 4-1 shows a typical visualization window.

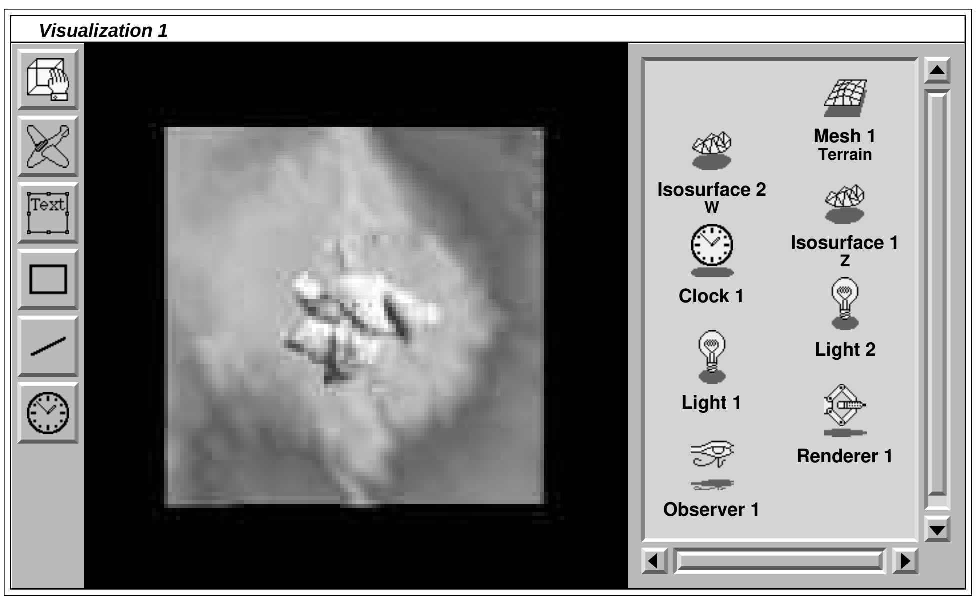

Figure 4-1. *Visualization window*

In the center is a 3-D **space**, which provides a world for the group of visualization objects in the window. There are two 2-D **space panels**, one in front of the space and one behind the space, which can be used to hold text and simple graphics. The front space panel is by default invisible, but the back space panel provides the black background of the space.

On the left is a column of icon buttons that select tools for interacting with the space or space panels using the mouse. Which tool is selected determines how the mouse behaves on the space. When a new visualization window is created, the space selection tool is automatically chosen. This tool allows you to move and rotate the space as well as select and edit objects on the space panels. Chapter 7 gives more information on how to use the other tools in the column.

On the right is an icon corral, which contains icons for all the visualization objects within the space. It also contains icons for **space controllers**, which control aspects of the visualization as a whole, such as lighting and point of view. New icons can be dragged into the icon corral, to replace or supplement objects already there.

You can hide the space tools and the icon corral by choosing Hide Control Panel from the Window menu. When the control panel is hidden, only the space appears in the visualization window.

Moving and rotating the space

The observer controller controls your position looking at the objects within the space. You can do basic motions of the observer around the space by using the mouse in the space itself.

When the space selection tool is selected, you can move and rotate all the visualization objects within the space using the mouse. Press and drag with the left mouse button to move the objects within the space. Hold down the Shift key while dragging to constrain motion to the nearest principal axis.

Note *When you move or rotate a space, you are really moving the observer within the space. You will not notice the difference from within the user interface, but you may need to know this when working on scripts for movies.*

Press and drag with the center mouse button to rotate the space. You can imagine that there is a virtual trackball covering the entire space. Where you click and drag determines the axis around which the space rotates.

Hold down the Shift key while dragging to constrain rotation to one of the principal axes of the space. If Rotation Inertia is set in the Preferences window (see Chapter 10), you can release the mouse button while the mouse is still moving to give the space a spin. To stop it spinning, simply click in the space.

If any lights are selected, the center mouse button will rotate the light sources rather than rotate the space. Rotation inertia does not work on light sources.

Hint *When a space is very complex or the computer is not very fast, interaction in a space may slow down. To improve speed during interaction, bring up the* Preferences *window and change the* Interactive Drawing *setting. For more information on preferences, see Chapter 10.*

Flight simulator

You can also move around in the space using a simple flight simulator, which is mainly for fun (so far, anyway). First make the visualization window as large as you can. Then select the icon button that looks like a plane.

When you press the center mouse button within the space, free flight will begin. The mouse works like a control stick. Push it forward to dive, and pull it back to climb. Move from side to side to turn. To accelerate, press the up arrow key. To decelerate, press the down arrow key. To stop forward motion and hover, press the zero (0) key.

Flight works best if the angle of view is wider than usual. (See "Observer" below.)

Space controls

Pull down Show Space Controls from the Windows menu to bring up the space control panel. The control panel has two check boxes. If Show Rotation Guides is on, spherical rotation guides will be shown in the space. If Show Motion Guides is on, a lattice of lines will be shown in the space. These can be helpful when moving and rotating in the space.

Space controllers

Space controllers, or **controllers** for short, are objects which control the way visualization objects are shown within a space. Like visualization objects, controllers have icons which appear in the visualization window corral. They also have control panels, which can be displayed by selecting the controller and choosing Show Controls from the Object menu.

Controllers can control more than one space. To make a controller control a second space, drag its icon to the corral of the visualization window containing the space, or choose Pick Up from the Object menu and click in the destination. In the case of some controllers, such as the Observer, Renderer, and Clock, there can only be one controller of each class controlling any one space; the new controller will replace the previous one.

Observer

The observer represents the user's point of view in the space. Observers have a location within the space, which is where the eye is located. They also have forward and up unit vectors which define their orientation and the distance to a **focus point**, which is where the eye is looking. Several other variables, such as angle of view and location of clipping planes, complete the description.

The observer can be moved and rotated by pressing and dragging within the space as described at the beginning of the chapter. All rotation occurs around the focus point. The other attributes of the observer can be controlled in the control panel, such as shown in Figure 4-2.

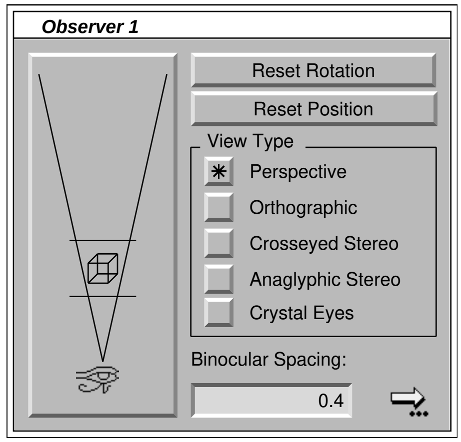

Figure 4-2. *Observer control panel*

At the left is a perspective control, which shows a schematic bird's-eye view of the observer (shown as a hieroglyphic eye) looking at the objects within the space. Sight lines radiate at angles from the eye giving the angle of view. Horizontal lines represent near and far clipping planes. The schematic cube represents the focus point of the observer, which is initially centered on the objects within the space.

To bring the objects in the space and the focus point closer or farther, press the cube in the control and drag it up and down. To widen or narrow the angle of view, press and drag the sight lines. To change the clipping planes, press and drag the horizontal lines up and down.

Hint *Press and hold down the Shift key to make fine adjustments to the control. This is especially useful with the clipping planes.*

When the view is Orthographic rather than Perspective, the sight lines appear as parallel vertical lines and specify the width rather than the angle of view. You can make the objects appear larger or smaller by narrowing or widening the width of view.

The Reset Rotation and Reset Position in the top right of the Observer control panel reset the observer to its original orientation and rotation. Below these are four radio buttons which determine the type of view displayed. When the Perspective button is checked, a single perspective view is shown. When the Orthographic button is checked, an orthographic view is shown. The perspective view shows objects more realistically. Objects closer to the observer appear larger than objects farther away. This is the best view for flying and is usually the best view for getting an overall picture of 3-D data. The orthographic view shows objects without foreshortening. Because objects are the same size no matter their distance from the observer, this view is useful for comparing portions of a field through a volume.

The other three buttons are for providing stereo 3-D views. When the Crosseyed Stereo button is checked, the two images of the stereo pair are shown side by side. To view them in 3-D, cross your eyes until the images converge, maintaining focus on the image. Some people find this very easy to do. Others find it takes a bit of practice. It may help to make the visualization window small or step back from the monitor. If viewing the image this way causes any eye strain at all, stop immediately.

When the Anaglyphic Stereo button is checked, the two images are displayed superimposed, one in red and one in cyan. To view the images, you need a pair of red/cyan 3-D glasses. This view is much easier to see than the cross-eyed stereo view, but because color is used to separate the images, the color of the resulting image will be washed out somewhat. Anaglyphic mode will only work if the Renderer is set to full color mode. (See "Renderer" below.)

When the Crystal Eyes button is down, stereo images will be displayed for use with the Crystal Eyes system. In order for this to work, you must have a set of Crystal Eyes goggles. Also, the visualization window must be set to full screen, and the control panel must be hidden. [Note: this is still under development and may change.]

The Binocular Spacing text box specifies the width between the eyes used for stereo calculations. This is the amount that the observer is translated to get the two images for a stereo view. The number is specified in space coordinates, in which the objects in the visualization are roughly 2 units wide and initially are 5 units from the observer. The default number given, 0.4, exaggerates true binocular viewing roughly twice, assuming that the image is viewed full screen at a distance of a meter or so. It provides a very good effect for large, smooth objects (such as the isosurface of the thunderstorm in the quick tour in Chapter 1). For more complex objects, this may be overwhelming, so it may be a good idea to crank it down a bit. At 0.2 it is very realistic. At 0.1 it begins to look flat, and at 0.0, of course, it is completely flat.

Stereo Color

Anaglyphic stereo filters the images through red and cyan filters before displaying. When choosing colors for an anaglyphic view it is best to pick shades of gray or colors with both red and cyan components, such as yellow and magenta. Imagine two vertical lines cutting the color wheel into three equally sized sections. Keep the colors in the middle third for the best effect.

The small arrow button in the lower right of the window is for accessing the actual parameters of the observer position and orientation. This is not needed for normal operation, but is provided for future use in creating complex moves in an animation. This will be discussed in the reference manual.

Lights

Lights are used to illuminate the objects in a space. They only work when the renderer is set to full color mode.

Lights work using the lighting capability provided by the graphics hardware. The number of lights that can be on at one time in a space is limited to the number of lights that the graphics system can handle. Most Silicon Graphics and IBM RISC System/6000 systems can handle at least eight lights. In practice, two are usually enough, though four can be useful for dramatic effects. All lights are located at infinity to increase speed and avoid inaccurate representation due to uneven lighting.

Items in the Surface control panel can affect how lights interact with objects. If two-sided lighting is checked in any visualization object's control panel, twice as many lights as specified may actually be used. See "Surface" in Chapter 5.

When the icons of one or more lights are selected in the visualization corral, lines showing the direction of the lights will appear within the space. The direction is always relative to the objects within the space. While lights are selected, you can rotate them to new positions by selecting the space selection tool and pressing and dragging in the space using the center mouse button. The process is the same as rotating the objects within the space. Hold down the Shift key while rotating to constrain rotation to azimuth or elevation alone.

Hint *When rotating lights, it may be handy to have the rotation guides displayed. Set the rotation guides using the space control panel, available via the* Show Space Controls *item in the* Window *menu.*

To turn lights on and off, select their icons and then choose Turn On or Turn Off from the Objects menu. The icons of lights which are off appear gray and dim.

The light control panel, shown in Figure 4-3, is very simple. It only contains a color wheel and a brightness control. To change the color, click on the desired location on the color wheel. Use the slider to change the brightness. Remember that the light color interacts with an object's base colors to produce the final colors of the object.

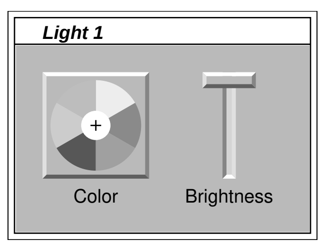

Figure 4-3. *Light control panel*

Renderer

The renderer controls the way objects within the space are rendered, or converted to images. The renderer control panel is shown in Figure 4-4.

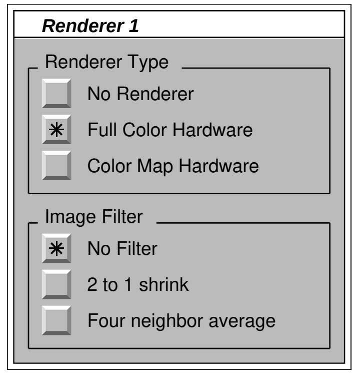

Figure 4-4. *Renderer control panel*

At the top is a group of radio buttons which specify the type of renderer. On most systems, there are two renderers: the Full Color Hardware renderer and the Color Map Hardware renderer. The full color renderer is most often used and fully supports all the shading and lighting features of the workstation. The color map renderer uses the color map hardware within the workstation. It does not do full light shading or transparency. However, it may be useful in 2-D visualizations where color is more important than shading. For more information on color modes, see Chapter 5.

At the bottom is a group of radio buttons which specify what filters to apply to the image after it is rendered. When the 2 to 1 shrink button is down, the image is shrunk to the bottom left quarter of the window, averaging groups of four pixels. This is useful with the Double Video Screen option in the Window menu for recording video without a scan converter. When the Four neighbor average button is down, each pixel is averaged with its four neighbors. The image filters only work when the renderer is set to full color mode.

Warning *Make sure no other windows overlap the visualization window when using the image filters. They read the pixels directly from the screen.*

Clock

The clock object controls the display of time-dependent data in a space. Since the concept of time is so important in SciAn, it is the subject of an entire chapter. Chapter 6 describes how to read time-dependent data and how to control and display time in visualization windows.

You can synchronize two spaces by dragging the clock icon from one into the icon corral of the other.

Creating visualizations

The easiest way to create a visualization for a dataset is to select the dataset within the Datasets window and press the Visualize button. SciAn will examine the dataset, pick a reasonable default visualization for the dataset, and open it in a new window. You can also drag a dataset icon into the corral of an existing visualization window to add its default visualization to the window.

A greater choice of visualization techniques can be obtained by pressing the Visualize As button. SciAn will bring up a window that looks like Figure 4-5.

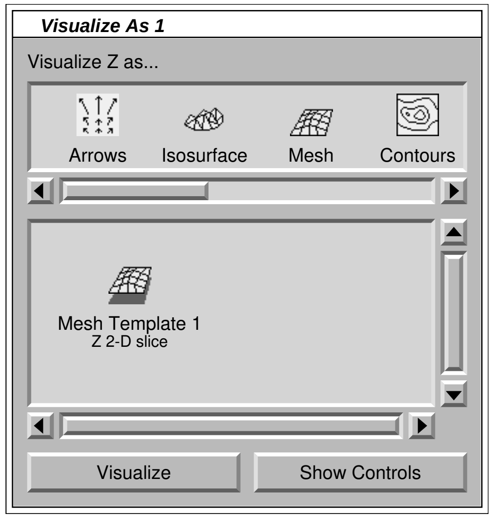

Figure 4-5. *Visualize As window*

At the top of the window is a horizontal series of icon buttons. Each button represents a visualization technique within SciAn. You may need to move the scroll bar below the buttons to see all the buttons in the field. Below the buttons is a corral containing ways of applying the selected visualization technique to the data. The contents will change depending on which button has been pressed.

SciAn may need to apply data modifiers to the datasets before they can be visualized using the chosen technique. The different icons in the corral represent different ways of modifying the data before visualizing it. For example, if you try to visualize a 3-D scalar field as a 2-D slice, SciAn will automatically apply an orthogonal slicer to the data. The kind of modification that has been done appears on the line below the name of the icon. If the 3-D scalar field were named Z, for example, this line would appear as Z 2-D slice.

To visualize any of the new objects, select them and press Visualize or drag them to existing visualization windows. Normally, when an icon is dragged from one window to another, the object will exist in both windows. However, the icons in the Visualize As window are actually **templates** for visualization objects, which work differently. When you drag them to other windows, new copies of the objects are created.

This is handy when making several slightly different visualization objects. Use Visualize As to bring up the visualization technique you want. Click on the icon and press Show Controls to set up the parameters. When they have been set up correctly, drag the icon to the window. Then change the controls and drag the icon to the window again. The two visualization objects will be set up with their own control panels initially set to be like the template.

Visualization objects

SciAn provides a variety of visualization objects for visualizing data. Each object can visualize one or more datasets. One dataset, called the main dataset, provides the basic shape and appearance of the visualization object. Other datasets may control aspects of the visualization, such as color.

All visualization objects have control panels, which you can show by selecting the icon for the object and choosing Show Controls from the Object menu. The control panel for every kind of visualization object has a column of icon buttons at the right. You may need to scroll down using the scroll bar to see all the icons. Each button brings up a different control panel, which controls a set of attributes of the object.

The button on the top brings up the top level control panel, which controls the most specific attributes of the particular visualization object and its main dataset. Once a visualization object has been created, its main dataset cannot be changed. Other control panels control more general attributes. These panels may have icon corrals for additional datasets, which may be changed. Each visualization object has a different set of control panels.

This section gives a brief description of the basic visualization objects within SciAn and gives the top level control panel for each. Chapter 5 gives more information on the possible other control panels each object may have.

Isosurface

The Isosurface visualization object is used to produce isosurfaces, or surfaces of constant value, from 3-D scalar fields. Isosurfaces can be made from structured scalar data fields defined over grids with three topological and three spatial dimensions. Isosurface can also be made from vector fields, in which case the absolute value of the vector is used.

Figure 4-6 shows the isosurface control panel. On the left is an icon corral containing the main dataset from which the isosurface is calculated. The slider in the center and editable readout on the left control the level at which the isosurface is calculated. At the lower left are two radio buttons which control whether the isosurface is intended to enclose values higher or lower than the isovalue.

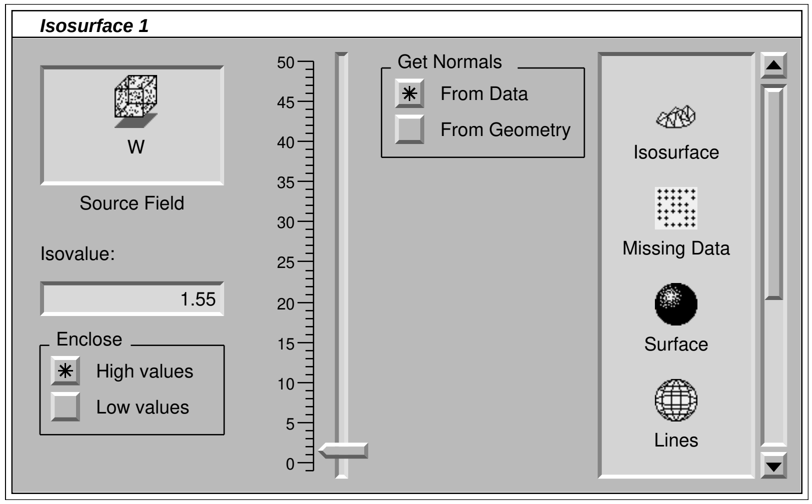

Figure 4-6. *Isosurface control panel*

Mesh

The Mesh visualization object is used to show a color mesh, or carpet plot of a 2-D scalar field, structured or unstructured. Figure 4-7 shows the Mesh control panel.

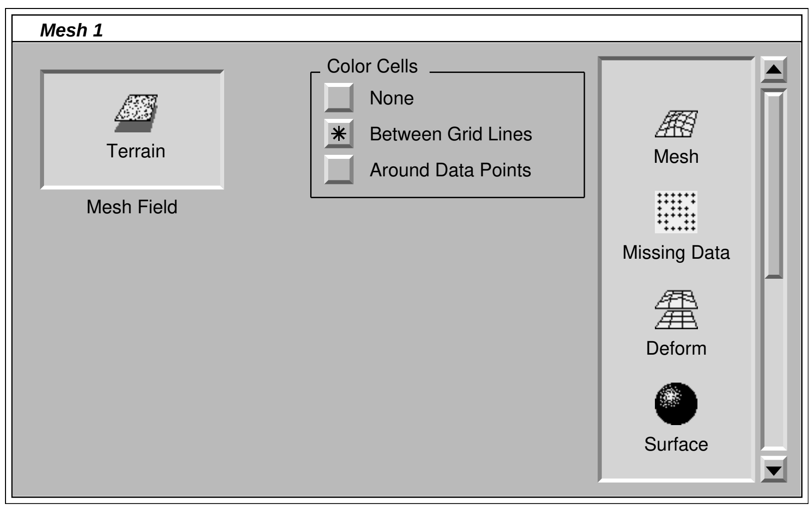

Figure 4-7. *Mesh control panel*

At the top left is an icon showing the main dataset. At the right is a series of radio buttons which determine how the color cells are formed. When the None button is down, no color cells are produced. Between Grid Lines is the most common setting. When this button is down, data points on the grid are used as vertices of polygons. When the Around Data Points button is down, polygons will be formed around the data points and the center of the polygons will be colored.

Note *How the color cells are formed interacts with how they are colored. When the color cells are between grid lines, it is most accurate to use smooth color shading. When the color cells are around data points, it is most accurate to use flat color shading. Other combinations will result in data interpolation to find vertex colors, which may hide data.*

The mesh object is a deformable object, so it has a Deform control panel as described in Chapter 5. Normally, the mesh is automatically deformed by the same field used to generate the cells and the colors, unless the deformation would be large compared to the size of the mesh.

See Chapter 5 to find out how to change the color field and shading.

Contours

The Contours visualization produces regularly spaced line contours of 2-D scalar datasets defined over structured grids. The Contours control panel is shown in Figure 4-8.

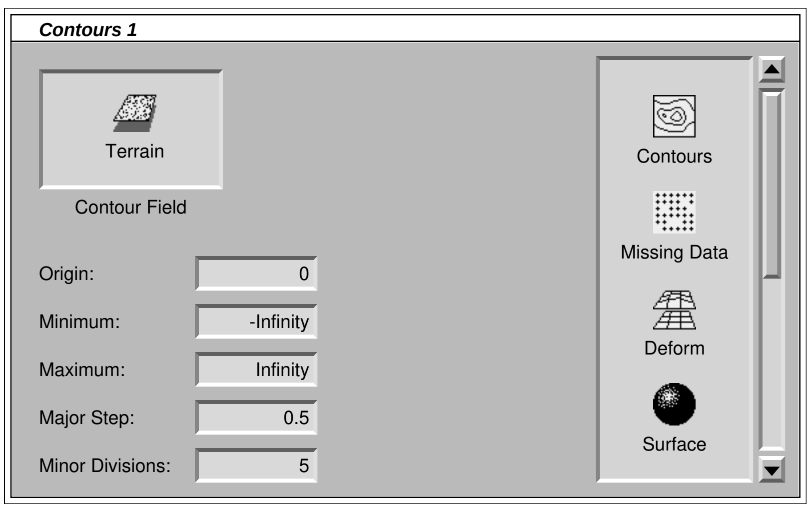

Figure 4-8. *Contours control panel*

At the top left is a corral containing the field used to generate the contours. The text boxes below the corral control the choice of contour lines.

Contours are calculated starting at an origin and moving up and down by a step value until the minimum or maximum is reached. The Origin text box specifies the origin, and the Minimum and Maximum text boxes specify the minimum and maximum contour values. Normally, they are set to –Infinity and Infinity.

The Major step text box gives the step size of the major contours. The Minor Divisions text box gives the number of minor divisions between each major division. For example, if Minor Divisions is 5, then one major contour will appear every 5 minor contours. The origin is always a major contour. Currently, major and minor contour lines look the same.

Hint *Contours are especially effective when used along with mesh objects. Try using different color schemes to color each object. You can deform contours using a scalar field just as you can deform meshes.*

See "Coloring visualization objects" in Chapter 5 to find out how to change the color field and shading.

Arrows

The Arrows visualization object is a simple way of visualizing vector fields with shaded arrows. Arrows can be applied to any vector field. Planes and lines of arrows can easily be shown using Visualize As. Each arrow is based on a data point and points in the direction of the field. Figure 4-9 shows the arrows control panel.

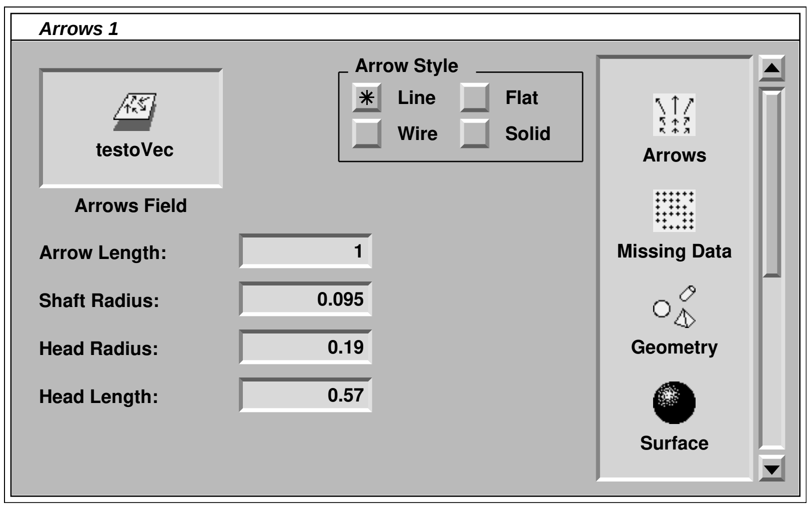

Figure 4-9. *Arrows control panel*

On the left is a corral containing the main dataset. Below this are four text boxes for entering numbers. The Arrow Length text box controls the length of the arrows. The magnitude of the vector is multiplied by this value to determine the length of the arrows. The Shaft Radius box controls the radius of the arrow shaft. The Head Radius and Head Length boxes control the shape of the arrowhead.

Hint *To get arrows that are shaped like pencils, make the head radius the same as the shaft radius. To get arrows that are shaped like long thin needles, make the head length infinite. Just type the word "infinity" into the text box.*

Point Cloud

The Point Cloud visualization object displays clouds of colored points, each point located at a data point. It is a simple and fast way of displaying an entire field at once.

A point cloud can be made of any scalar dataset defined over a structured or unstructured grid. Figure 4-10 shows the Point Cloud control panel.

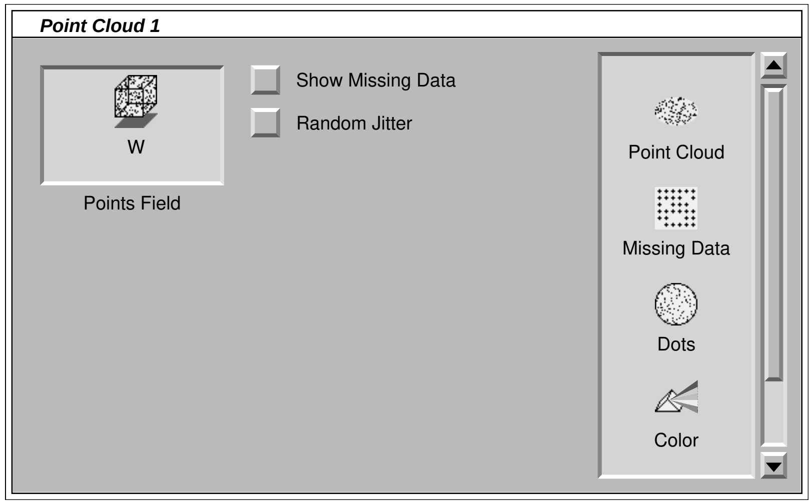

Figure 4-10. *Point Cloud control panel*

The Show Missing Data check box at the upper left controls whether missing data points are shown. If checked, missing data points are shown using the missing data color, which initially is a dull gray. If it is not checked, missing data points are not shown at all.

The Random Jitter check box controls whether random jittering is applied to the points. Normally, points are drawn exactly where they are located within the grid. This is precise, but the regularity of the grid usually results in an excessively regular appearance which may distract from the information in the dataset. If this check box is checked, each data point is randomly jittered within a box the size of a data cell around the data point. This results in a smoother appearance which may be better for getting an impression of what is happening within the field.

Geometry

The Geometry visualization object displays geometric objects directly as read from a file. Geometric objects are made of lines, polygons, spheres, and conical frusta.

Geometric objects can be colored by a fixed color, by scalar fields, or by themselves. By default, they are colored by themselves. The colors of individual geometric primitives can be controlled by bringing up the color palette of the object. All the base colors used in primitives are located in the color palette.

See Chapter 5 for more information about coloring objects and about the way geometric objects are displayed. Also see the NFF file reader in the reference manual, Chapter 4, "File Formats."

Numbers

The Numbers visualization object produces a numeric display of a single scalar value. The numeric display appears on the front space panel and behaves very much like a clock display or annotation (see Chapter 7). Figure 4-11 shows the control panel of a numbers object.

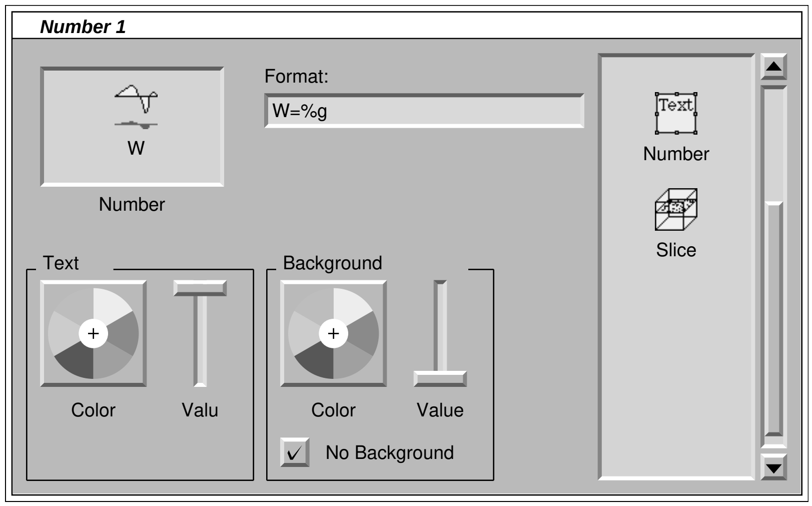

Figure 4-11. *Numbers control panel*

At the top right is a format string which is used to print the number. It works very much like the format for `printf` in the C programming language. [Kernighan and Ritchie, 1978.] For more information on the format, see the discussion on time formats in Chapter 7.

At the bottom of the panel are two groups of controls to set the text and background color. By default there is no background, so you can see through the text box where the text does not appear. Uncheck the No Background check box to change this.

The text font, size, and alignment can be changed by selecting the text box within the space and using items in the Text menu. For more information, see Chapter 7.