

Chapter Four: Related Work in the Visualization of Multiple Spatial Variables

In this Chapter I describe several examples of past and current work in the field of Scientific Data Visualization. Each technique bears similarities to DDS, in either the type of data displayed, the graphical methods used, or in the number of variables the technique can show.

Ware on page 188 of his book, *Information Visualization, Perception for Design* [2000] describes the class of Glyph data visualizations as either *integral* or *separable*. Both types of display map data variables to two or more visual attributes of a single display object, or Glyph. An *integral* display maps data values to attributes that are perceived holistically, such as the length and width of an ellipse. A *separable* display maps data to graphical attributes that are perceived separately, such as the color and orientation of an ellipse.

In the descriptions below, the data visualization techniques are classified as either *layered* or *integrated*, where I extend Ware's definition of *integral* to mean that individual variables contribute to an integrated display and are not visually separable in the final image. In *layered* techniques each data value is visually separable from the others. Often techniques are both layered and integrated, for example if the images contain multiple separable layers, but within the layers integral display elements are used. Figure 4.3 is an example of a layered technique that uses integral display elements from Kirby, Maramanis, and Laidlaw [1999]. It uses multiple semi-transparent layers to display fluid flow. DDS, on the other hand, is a layered, separable, display technique.

I first describe data visualization techniques that are layered and separable, then layered techniques with integral components, and finally describe some examples of integrated methods. For the most part, this classification is my own; however, where possible I refer to work done by the authors related to visual separation of the displayed data. Table 4.1 provides a summary of the work covered in this chapter.

Authors	Year	Type	Figure	Similarity to DDS
Crawfis and Allison	1991	Layered	4.1	<ul style="list-style-type: none"> • Functions in separate layers
Turk and Banks	1996	Streamlines	4.2	<ul style="list-style-type: none"> • Separate graphical layers
Kirby et al.	1999	Painterly Rendering	4.3	<ul style="list-style-type: none"> • Graphical layers • Transparency
Laidlaw et al.	1998	Painterly Rendering	4.4	<ul style="list-style-type: none"> • Graphical layers
Healey and Enns	2002	Painterly Visualization	4.5	<ul style="list-style-type: none"> • Same data type
Healey et al.	2001	Pexels	4.6	<ul style="list-style-type: none"> • Similar data type • Perceptual separation of data
Weigle et al.	2000	Slivers	4.7	<ul style="list-style-type: none"> • Same data type • Separate Graphical layers
Smith et al.	1991	Icon/Glyph	4.8	<ul style="list-style-type: none"> • Similar data types

Table 4.1: A summary of the related work discussed in this Chapter, in the order discussed.

Layered Visualization – Crawfis and Allison

Crawfis and Allison [1991] describe a multivariate visualization solution, which displays both scalar and vector data with up to four different visual layers (Figure 4.1). Each visual layer uses a different display feature, such as an underlying color image, a surface grid, contour lines, and a vector plot. Crawfis and Allison use bump maps to display the contour lines and the vector field, because they observed this scheme reduced clutter in the image, and interfered less with the underlying color. The authors' observation agrees with the experimental evaluation of DDS, which showed that bump-map layers do not interfere with the alpha-blended targets. However bump-map layers interfere with each other, and the number of bump-map layers should be limited to one or two, as in the work of Crawfis and Allison. Although the authors do not report psychophysical studies related to their technique, each layer appears visually separable: the bump map clearly stands out from the underlying color map and grid lines.

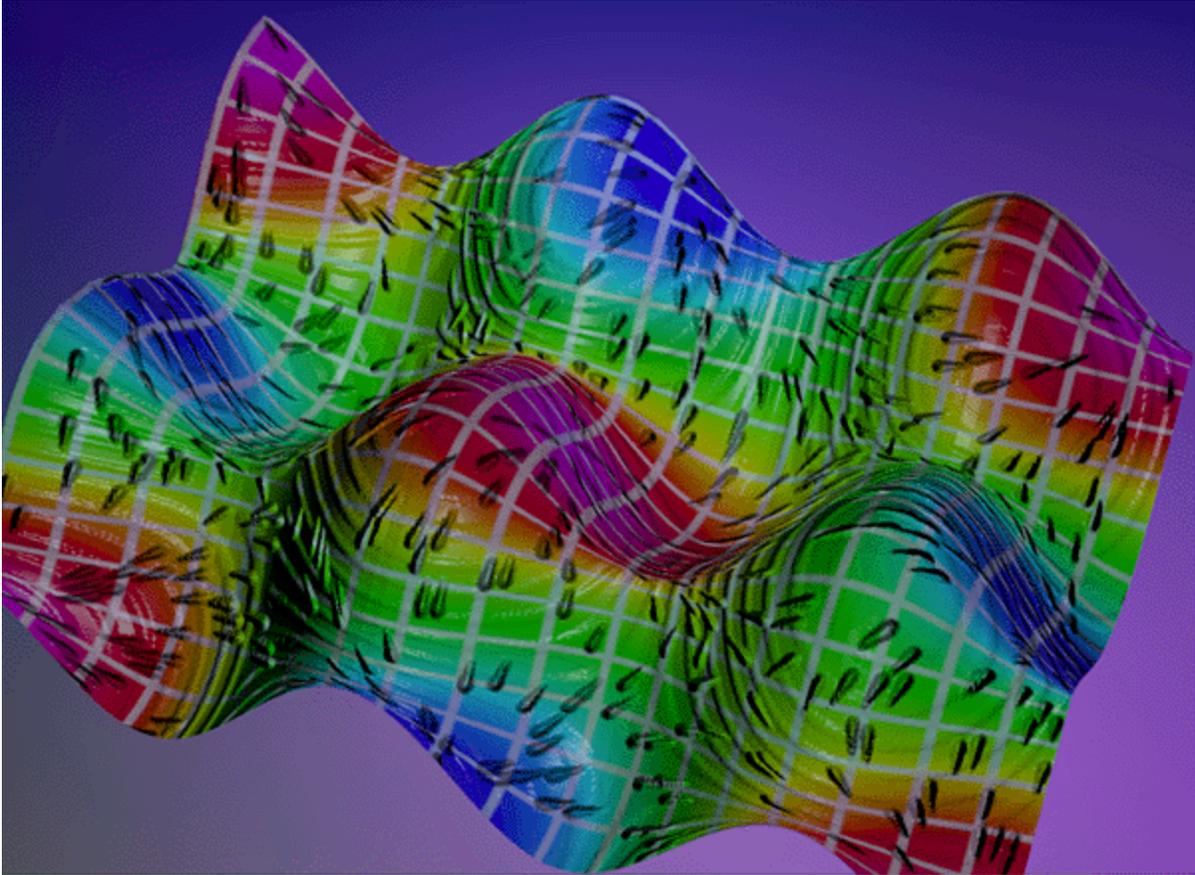


Figure 4.1: The image shows five artificial functions displayed with the surface shape, surface color, raised contour lines, bump-mapped vector field, and grid. Image courtesy of Crawfis [1991].

Streamlines

Streamlines, by Turk and Banks, [1996] is a technique for displaying flow information. It is often useful when combined with an underlying color image of the geographical terrain, as shown in Figure 4.2, and can therefore be considered a layered visualization. Each streamline displays the flow direction vector and the flow magnitude. Arrowheads indicate the flow direction; thickness, density, or intensity of the lines indicates flow magnitude; icon curvature shows vorticity. In [Ware, 2000] orientation and line width are considered moderately integrated, less integrated than line width and line length but more so than size and color.

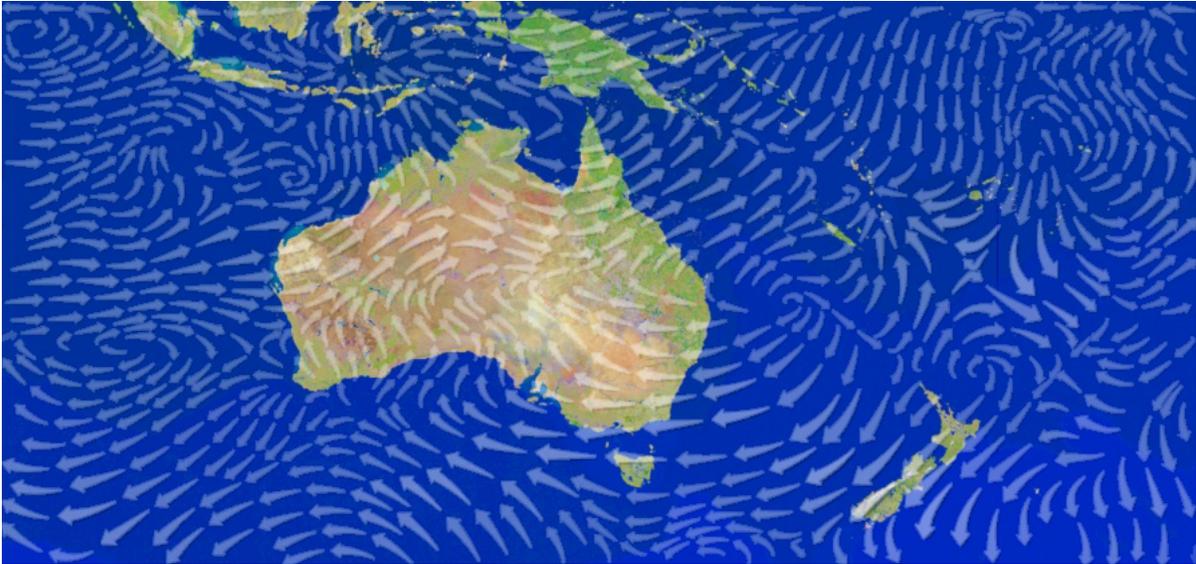


Figure 4.2: The image shows wind patterns over Australia: wind direction is indicated by the direction of the arrows, wind speed is indicated by the size of the arrows. Image courtesy of Turk [1996].

An interesting aspect of the Streamlines display algorithm is that it uses an iterative process to change the positions and lengths of the streamlines, to join streamlines that nearly abut, and to create new streamlines to fill gaps, with the goal that the streamline placement is neither too crowded visually, nor too sparse. This self-assembly-like process of generating the streamlines is also used to distribute the oriented lines in the Slivers technique described below. It also holds promise for DDS, both in terms of distributing the Gaussians within a sampling array, and in terms of adjusting the overlap between layers so that the different Gaussian layers with similar sizes have little overlap.

Painterly Rendering

Kirby et al. [1999] use a *Painterly Rendering* technique, described in [Laidlaw et al. 1998], to illustrate fluid flow. The technique is inspired by the brush strokes used in oil painting. The authors relate that they borrow rules from art to guide their choices of colors, textures, and the visual display elements.

Their technique is used to visualize four different components of fluid flow: the direction and speed of the flow, the vorticity, which represents the rotational component of the flow, and the rate-of-strain tensor of the flow, which is the rate at which fluid elements change shape as they move through the flow field. They explain the rate-of-strain tensor for fluid flow as small circular fluid elements compressing into ellipses as they move through the flow field [Kirby et al. 1999].

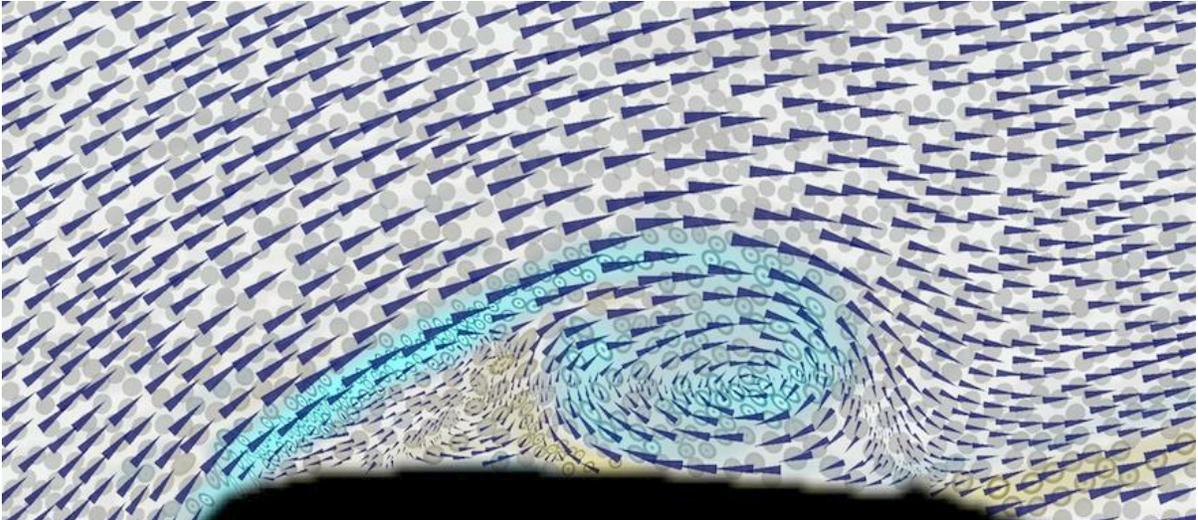


Figure 4.3: The image shows measured fluid flow: the velocity is shown with arrow direction, speed is shown with arrow area, vorticity with the underpainting color and ellipse texture, the rate of strain is shown with the ellipse radii, divergence with ellipse area, and shear with ellipse eccentricity. Image courtesy of Laidlaw [1999].

Figure 4.3 from [Kirby et al. 1999] shows an example. In the image there are five layers. First, a bottom layer of light gray, which is chosen to contrast with the colors in the upper layers. Second, the underpainting, illustrates vorticity with semi-transparent color, blue showing clockwise rotation and yellow showing counter-clockwise rotation. The magnitude of the rotation determines the level of transparency, with higher vorticity represented by greater opacity – this is similar to how DDS uses transparency. The third layer is the ellipse layer, which shows the rate-of-strain tensor – how the fluid elements compress as they move through the flow. Uncompressed elements show as circles, compressed elements show as ellipses where the orientation of the ellipse represents the flow direction and the area of the ellipse represents flow divergence. Ellipses are textured to illustrate vorticity magnitude; higher magnitude vorticity is displayed with a textured dot in the center of the ellipses. This is redundant display with the underpainting opacity. The fourth layer is a layer of arrows to show the flow direction and magnitude. The authors remark that the data, which was measured from passing an airfoil through stationary water, shows compression or out of plane motion of the fluid, which explains why the ellipses do not have constant area as is expected with incompressible flow [Kirby et al., 1999].

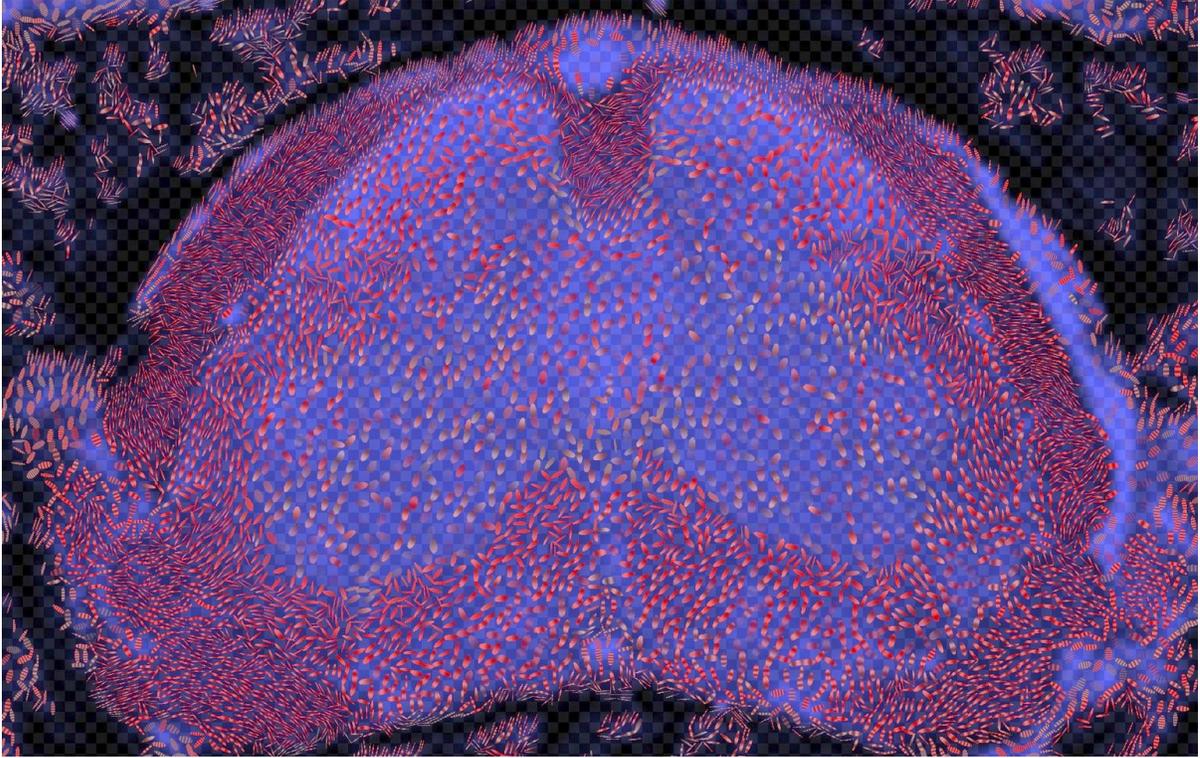


Figure 4.4: The image shows MRI diffusion rate images from the spinal cord of a healthy mouse. The anatomical image is depicted with the lightness of the underpainting, the voxel size is shown with a semi-transparent checkerboard, the ratio of largest to smallest eigenvalue is shown with the length/width ratio of the ellipses and with the ellipse transparency, the principal directions of diffusion in the x,y plane are shown with the ellipse direction, the principal direction of diffusion in the z direction is depicted with red saturation, and the magnitude of the diffusion rate is shown with the ellipse texture frequency. Image courtesy of Laidlaw [1998].

Within the arrow and ellipse layers, each icon carries two or three different data variables. For the arrows, orientation, and area are used to display velocity and magnitude, as in the Streamlines technique. For the ellipses, the shape and texture carry information; shape shows compression and texture shows vorticity magnitude. For both the arrow and ellipse, display dimensions are integral.

When I look at the image in Figure 4.3, first the arrow layer and the color underpainting stand out, while the ellipse layer requires more focused attention. This matches the authors' goal in visualizing the data, which is to guide the viewer to a better intuition of fluid mechanics, with the viewer first understanding velocity, then vorticity, and finally rate-of-strain tensor and how all relate.

The *Painterly Rendering* technique was first introduced to display diffusion-rate images by Laidlaw et al. [1998]. Diffusion-rate images, which are calculated from a set of magnetic resonance images, show the rate at which water molecules diffuse through tissues. Because water molecules move at different rates in different tissues – faster in liquids and slower in muscle – diffusion rate images segment the data by tissue type. The authors demonstrate visualizations of diffusion rate images for both a healthy and diseased mouse spinal cord, with the goal of aiding diagnosis.

Figure 4.4 from [Laidlaw et al. 1998] shows their work. In the figure there are three layers. The first layer is the underpainting; it shows the anatomical image associated with the diffusion rate image. The second layer is a semi-transparent checkerboard, which displays the voxel size of the data elements. The third layer is the brush stroke layer, which displays information about the diffusion: the stroke direction and amount of red saturation show the rate of diffusion, and the stroke texture frequency shows the magnitude of the diffusion. Stroke length represents diffusion isotropy; for areas in the image where there is anisotropic diffusion, typically in white matter, the strokes are elongated, whereas in areas where there is isotropic diffusion, typically gray matter, the strokes are spherical [Laidlaw et al. 1998].

The images resemble paintings more than traditional data visualization images, and indeed the authors report that they derived much of their inspiration from Van Gogh:

"We found Van Gogh's style particularly applicable because he used expressive, discrete strokes that are combined to represent a continuous scene. They can be read differently from different distances and can encode different information from each of those viewpoints."
(page 129).

Although DDS was not directly inspired by the work of artists, it has a similar flavor. Much of the work described later in this chapter was derived from perceptual psychology, which describes the basic visual features that are easily seen by people. DDS, like the work of Laidlaw et al., was developed with both esthetics and clarity of display in mind. In addition, DDS images at different viewpoints, from arm's length or on close inspection, reveal different characteristics of the data. One participant in the pilot study described in Chapter Three remarked on this explicitly, saying that the overall shapes "jumped out" at him when he leaned back and looked at the images from greater distance.

Painterly Visualization

The search for art in data visualization inspires the work of Healey and Enns [2002], who are attempting a more rigorous investigation into how people perceive art and data. They compare aspects of impressionist painting to important cues in visual perception: color and lighting they relate to hue and luminance, whereas brush stroke styles such as the stroke path, density, and length they relate to orientation, contrast, and size. Data attributes map to specific painterly styles by varying the visual appearance of the brush strokes with the underlying data.

Figure 4.5 shows an example from [Healey, 2001] of climate data similar to the data displayed in Chapters One and Two. In the image, temperature maps to the paint color (lightness), pressure maps to brush stroke size (area), wind speed maps to stroke density (calm winds produce less dense application of strokes), precipitation maps to stroke orientation (heavy rain produces more slanted strokes), and frost frequency maps to stroke regularity (heavy frost produces a less regular placements). Each display element is a combination of integrated (size, orientation), separable (color), and group (density, regularity) characteristics. When I look at the image in Figure 4.5, color stands out first, followed by density, orientation and size. In denser and darker areas orientation, size, and regularity are more difficult to see.

My own impression closely matches psychophysical studies the authors describe and in earlier work, where color is identified rapidly, with little interference from variation in the other dimensions, and where regularity is harder to identify [Healey and Enns, 1998; 1999].

In addition to investigating the relative salience of the data mappings in the artistic renderings, Healey and Enns investigate the esthetic appeal of the generated images. They found that people who preferred abstract images to begin with rated the painterly visualizations most artistic out of a selection of computer-generated non-photorealistic rendered images, and well-known abstract and impressionist works.

Healey's work in *Painterly Visualization* has furthered my own interest in how DDS can be used to display data in both an informative and esthetically appealing manner. The DDS visualizations of climate data, presented in Chapter Two, show very different patterns in the data for the months of February and August. Visually appealing visualizations, because they engage our interest and careful study, can improve our intuition of the changes in weather across the seasons – even those for which the individual layers are not visually distinguishable.

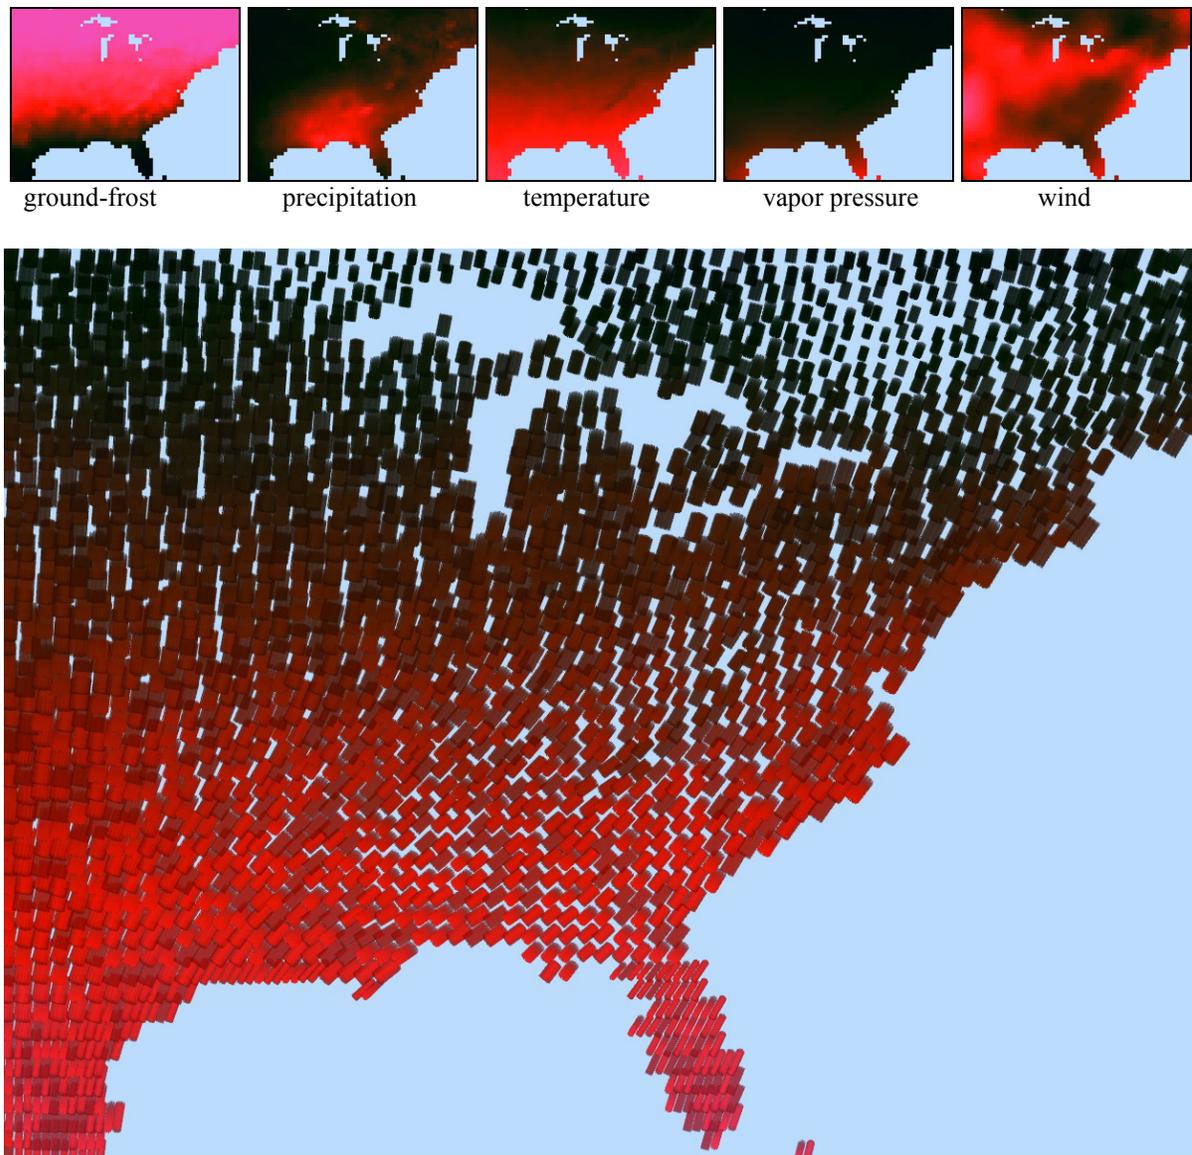


Figure 4.5: Depicted are climate data from the Intergovernmental Panel on Climate Change. Temperature is shown with a color gradient, pressure with stroke frequency, wind speed with stroke coverage, precipitation with stroke orientation, and ground-frost frequency with stroke regularity. Image courtesy of Healey [2002].

Pexels

The painterly visualizations are similar to earlier work by Healey et al [1995, 1998], which investigated the use of *Perceptual texture elements*, or Pexels, for displaying data. Pexels are similar to the brush strokes used in the painterly visualizations in that a single pexel may vary in height and color, while pexel placement varies in density and regularity, depending on the underlying data to be displayed.

Figure 4.6 shows an example of Pexels used to display Typhoon Amber as it approaches Taiwan and makes landfall, where wind speed maps to height, pressure maps to density, and precipitation maps to the regularity of pexel placement on the grid.

In the experiments reported in [Healey and Enns, 1998; 1999] the authors investigate the viewer's ability to identify rapidly and accurately target pexels in an array of background pexels. They tested two pairings of height differences of the target pexels compared to the background pexels; a medium-short target-background pairing and tall-medium pairing. They tested three different exposure durations: 50, 150, and 450 milliseconds. Density and regularity were also varied in background and target as well as the target patch size, which was either 2x2 or 4x4 pexels.

The results showed that some dimensions were more salient than others. Tall pexels stood out from medium pexels with little interference from random density and regularity changes in the background. However, shorter, denser, and sparser targets were more difficult to identify and suffered from interference effects due to variations in height, regularity, and density of the background pexels. Regularity was not well identified. When color was added to the pexels it was detected rapidly with little interference from variations in the background pexel dimensions. However, variations in background color did interfere with the identification of pexels distinguished by height and/or density.

Healey's work in Pexels was a strong influence for DDS layers to be perceptually separable – that viewers should be able to study the contributions of each variable layer with minimal interference from the other layers displayed.

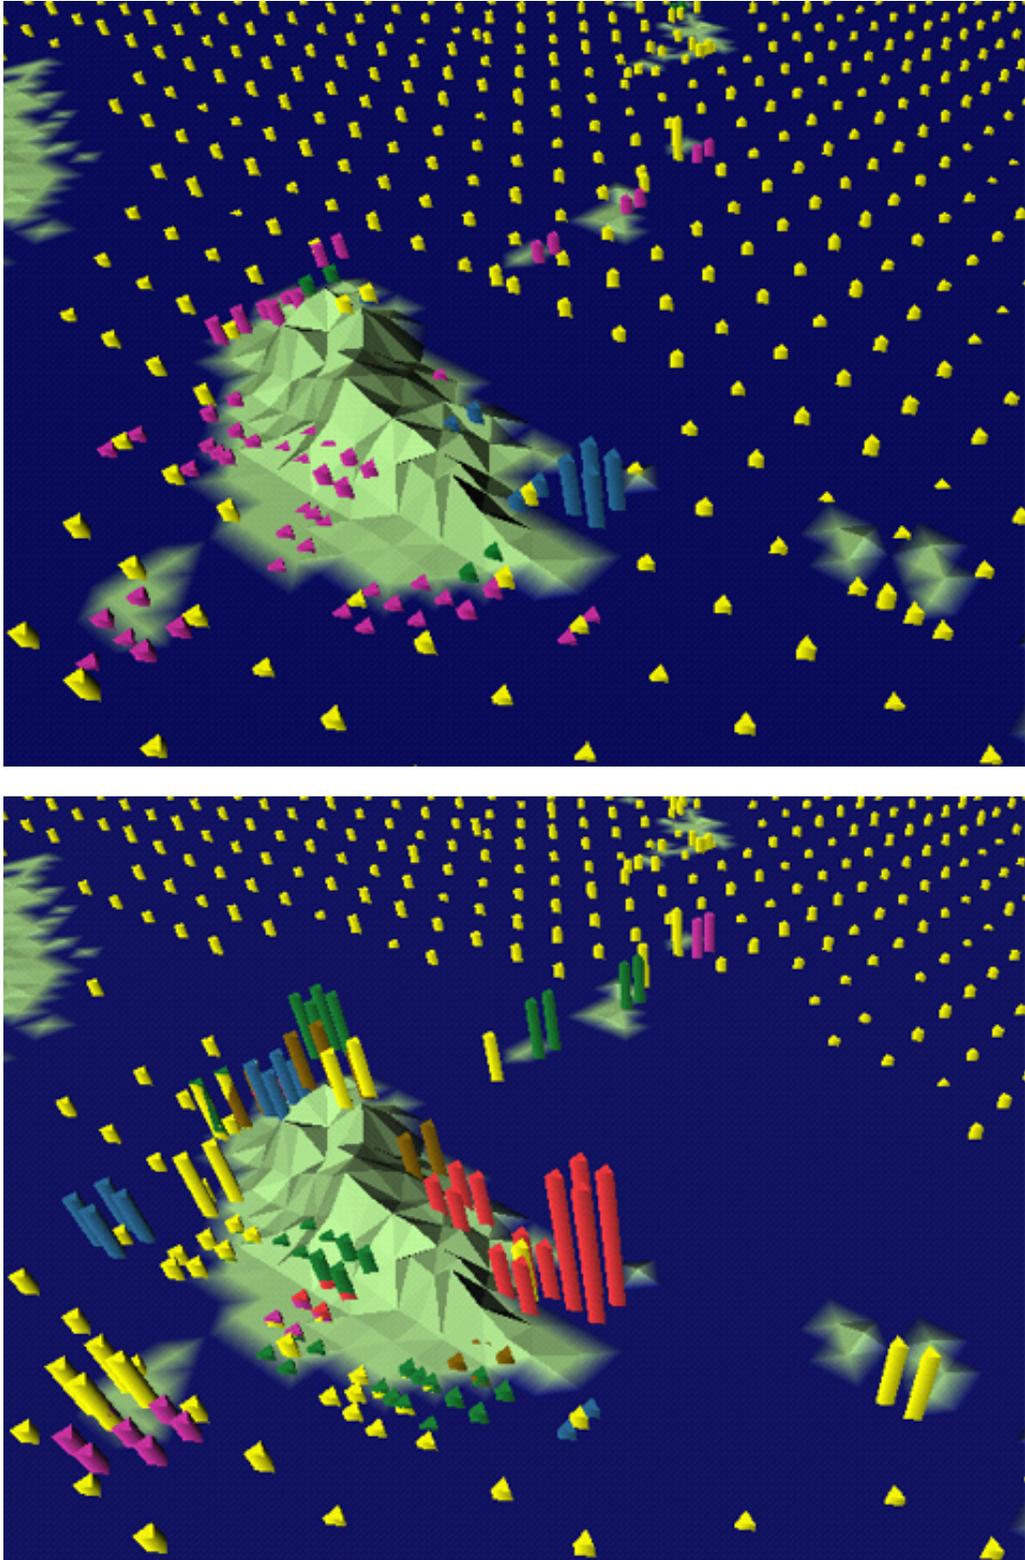


Figure 4.6: The images show typhoon Amber over Taiwan at two points in time as the typhoon makes landfall. Pressure is displayed with pexel density, precipitation with pexel color, and wind speed with pexel height. Images courtesy of Healey [2001].

Slivers

The *Slivers* technique, by Weigle et al. [2000], uses orientation and luminance to display up to eight or nine layers of scalar data. *Slivers* was developed in the same lab at the same time as DDS, with the idea that it would display the same data but perhaps emphasize different characteristics of that data. In the paper the authors describe the goal of the *Slivers* technique as *the rapid visual exploration of multiple overlapping scalar fields*, and they claim *Slivers* images allow viewers to discriminate data values in individual fields while at the same time highlighting the interactions among fields [Weigle et al. 2000].

Figure 4.7 shows an example of the *Slivers* technique applied to the same SEM data set used in Chapter Two. Each variable layer is displayed with a sliver texture with unique orientation. The luminance of the sliver is directly determined by the data value at points where the sliver is displayed using a perceptually-balanced luminance scale. *Slivers* are blended based on a maximum-value function, so that at points where slivers cross, the sliver showing the data point with the highest value is displayed.

The different *Sliver* orientations were chosen based on a psychophysical study. A target patch of slivers was tested against 19 different background orientations, where the orientation of the target differed from the background in varying amounts. They found that an orientation difference between the target and background *Slivers* of 15 degrees was enough for the target patch to be rapidly and accurately perceived.

The authors claim that thin, well-separated slivers is key to allowing values from multiple fields to show through. Moreover when two *Slivers* textures overlap, they combine to form a new recognizable patterns such as plus, cross, or star-shaped. For this reason, the authors describe the *Slivers* technique as being closely related to the *Glyph* techniques. The emergent features in the *Slivers* technique fall under Ware's definition of an integrated display [Ware, 2000] in that once two sliver layers combine to form a new shape they are no longer perceived individually.

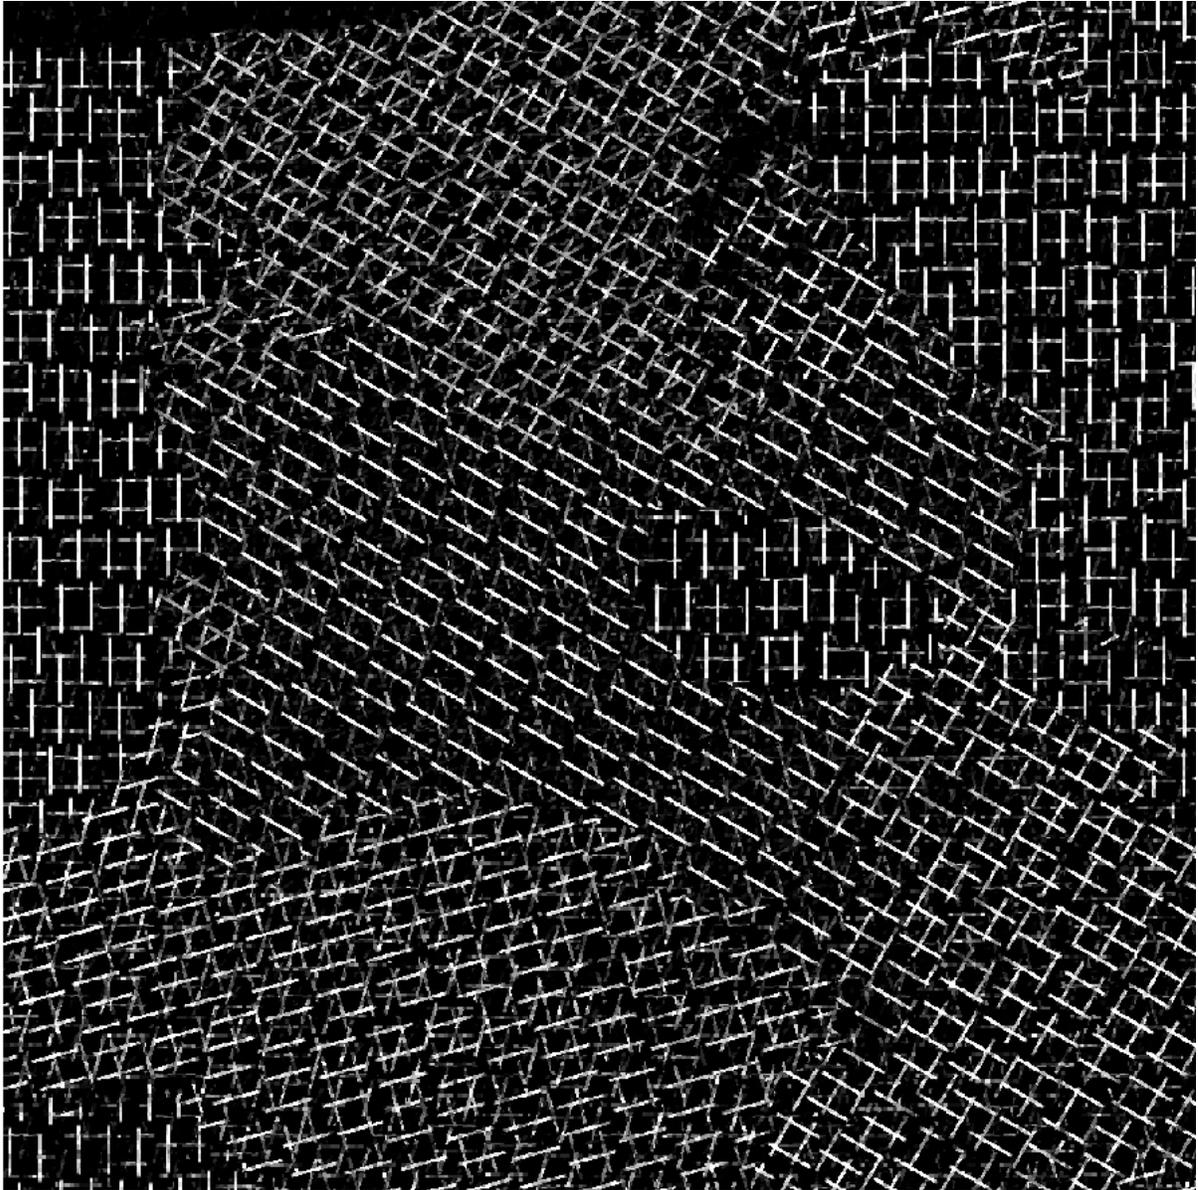


Figure 4.7: The image shows the same SEM data shown in Chapter Two of this dissertation. Calcium is shown with a 15° rotation, copper with 30°, iron with 60°, magnesium with 75°, manganese with 105°, oxygen with 180°, sulfur with 150°, and silicon with 90°. Image courtesy of Weigle [2000].

Three distinct areas of texture stand out in the Slivers image in Figure 4.7: the texture dominated by the Slivers of vertical orientation at the left and right edges of the image, the texture dominated by Slivers with north-west to south-east orientation, found in the center of the image, and the texture with strong south-west to north-east orientation, found in the bottom and upper areas of the image. In the psychophysical study presented in the paper, the authors do not describe a test of the number of orientations and therefore the number of discriminable variables viewers can see in a multi-textured Slivers image. This is important, as the visual salience of an object decreases with the variety of distractors, regardless of preattentive visual attributes among individual layers [Ware, 2000]. The number of easily distinguishable Slivers may be fewer than the eight shown, and Ware argues that perhaps only three orientations are easily distinguished in a Glyph-type texture display [2000 p. 197].

Chapter Two presents DDS visualizations of the same SEM data shown in Figure 4.7. In DDS images the individual layers are perhaps more visually distinct than in the Slivers images, however, one could argue that the combination of sliver layers forms more recognizable patterns where data is spatially correlated than the combination of DDS layers. However, the distinct regions shown by the “plus” pattern in the slivers image in Figure 4.7, are also distinct combinations of colored spots in the DDS images of the same data.

Icons and Glyphs

The Exvis system described in [Pickett and Grinstein, 1988; Grinstein, Pickett, and Williams, 1989; and Smith, Grinstein, and Bergeron, 1991] creates visual textures to represent high-dimensional data (up to 30 variables). The textures are based on densely packed icons, where each individual icon is a stick figure that encodes data in the orientations, lengths and positions of the limbs. They describe looking at wall to wall carpet as an analogy to their work – foot prints and brush patterns caused by the vacuum cleaner create visually different areas in the carpet. The stick figure they liken to the individual fibers of the carpet where the changes in size, shape and spacing of the icons creates visually different areas [Grinstein et al. 1989].

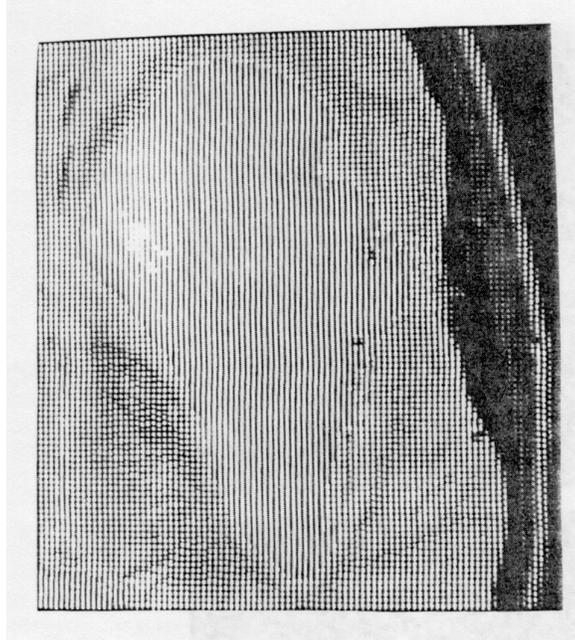


Figure 4.8: The images shows two MRI cross sections of a brain with a tumor, the tumor hot spots is visible in white toward the upper left of the image Image courtesy of Smith [1991].

The authors derive the motivation for the stick figure icon from visual perception literature, which has shown line orientation to be a strong visual cue. Icons have both a body and four limbs, where limbs can be attached only to the ends of the body or to the ends of other limbs, with the limit that only two attachments can be made at any point. Data values map to the length, width, orientation and color of the limbs. Thus, a five limbed icon can represent up to 20 spatial variables [Grinstein et al. 1989], many more than have been explored with DDS. I believe that if 20 variables were displayed with DDS, the result would be vary similar in nature to the Glyph techniques, in that distinct patterns in the final image would be visible, but that no one individual variable could be visually distinguished from the rest.

Although 20 spatial variables can be displayed with the Exvis stick icon, the data values are essentially input parameters to a texture generation algorithm. Once the texture is generated, it is no longer possible to distinguish visually which of the 20 variables is high or low in any given location. While this is not the goal of Exvis, it is the goal of DDS. Figure 4.8, from [Smith et al. 1991], shows an example from a composite of two MRI brain scans.

Both the addition of sound [Grinstein et al. 1989] and the importance of real-time interaction with icon limb configurations and data-to-parameter mapping [Smith et al. 1991] have been explored with the system. Levkowitz [1991] chose to use a different icon configuration, based on a pie chart, where the color of each icon limb changes with the data while icon shape remains constant. Adding variables to the display adds limbs to the icon, for up to six variables.

Chernoff Faces [1973] is a very different example of an integrated, or glyph, technique, less related to DDS. This technique is used to display up to 18 different variables of non-spatial data. The variables map to features of a cartoon face, such as the length and curvature of the mouth, the size of the eyes, position of the pupils, and the size and orientation of the eyebrows. The author illustrates this technique on a set of 87 geological specimens from a limestone formation. The goal of the visualization was to enable the researcher to group the specimens based on measurement values. A face was created for each specimen, and the researcher could then group the faces based on their visual characteristics. This technique allows the researcher to quickly understand relationships and patterns in the data, based on the relevant information used to group the faces. This aids the researcher in that it narrows and focuses the choice of analysis.

Mapping data to facial features is a fascinating idea, as a large part of our visual system is devoted to face recognition. Humans have the ability to recognize and characterize a large range of emotions based on facial expression. For example, a Chernoff face with large eyes can express several different emotions -- from surprise, to anger, or friendliness -- depending on how the mouth and eyebrows are drawn. It is extraordinary how the cartoon faces change expression with data values.