

Multivariate Visualization With Attribute Block Displays

James R. Miller
University of Kansas
EECS – 2001 Eaton Hall
1520 W. 15 Street
Lawrence, KS 66045-7621 USA
jrmiller@ku.edu

ABSTRACT

Visualizing multiple continuously-varying attributes across a region of space is a complex problem. We have recently introduced a visualization scheme based on Attribute Blocks, a regular array of cells, each of which is assigned an attribute. The Attribute Block array can be dynamically sized and tiled across a field of view to allow an arbitrary number of attributes to be visualized at once on a single display. The technique imposes no limits on the number of attributes that can be visualized at once, however human perceptual limitations clearly do. In this paper we briefly discuss related literature, review the Attribute Block method, and then discuss ongoing work whose goal it is to develop strategies for designing Attribute Block displays which maximize users' abilities to visualize how multiple attributes vary across a region of interest.

Keywords

Visualization; Multidimensional multivariate visualization

1. INTRODUCTION

Multidimensional multivariate visualization refers to the visualization of n attributes as a function of their position in d -dimensional space [Won97a; deO03a; Lov05a]. A variety of visualization and exploration tools have been investigated, both for navigating the d spatial dimensions as well as for presenting how the n attributes vary across the d -dimensional spatial domain.

We recently developed a tool for the visualization of n attributes continuously defined across a geospatial region (i.e., $d = 2$) [Mil07a]. We dubbed the tool *Attribute Blocks* since the display was based on a regular array of Attribute Block *Cells*, one cell being used to visualize the variation of an attribute across the small region covered by the cell. Initial experiments indicated that several interesting areas of the data set were very rapidly detected as preattentive features. However it was also clear that considerable

additional work was required to optimize the displays for specific applications. For example, selection of appropriate color ramps for attributes was vitally important.

Moreover, while our initial approach assigned each (i,j) position inside an Attribute Block array to a unique attribute, we discovered that our ability to discern patterns in the data could be greatly enhanced by generating certain structures in the Attribute Block array by repeating attributes in multiple positions. We have explored some aspects of this dimension of the use of Attribute Blocks and present our preliminary results below.

In this paper we first review several relevant techniques in the area of multidimensional and/or multivariate visualization. Then we turn to the use of Attribute Block displays in the $d = 2$ case and show how they can be used to enhance our understanding of how multiple attributes (as well as quantities derived from them) vary across a region.

2. PREVIOUS WORK

The quantities we wish to visualize correspond either to discrete objects or events whose location is one of the attributes that must be visualized or to continuous functions that are defined everywhere. Examples of the former might include sightings of endangered species, voting results, or locations of windmills.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright UNION Agency – Science Press, Plzen, Czech Republic.

These are typically represented using icons displayed at the points, oftentimes with additional attributes mapped to the type, color, shape, topology, etc. of the icon.

Examples of continuous attribute functions might include average temperature or height above sea level. Frequently continuous phenomena like these might only be *measured* at discrete locations, but clearly they have values everywhere, hence we view them as continuous even when they are only captured at discrete locations.

Our focus in this paper is on the visualization of continuously defined attributes. While some techniques are applicable to both domains, others (e.g., pixel maps [Kei04a] or traditional scatter plots and multivariate extensions including trellis plots) are uniquely suited to the discrete domain. Moreover, the focus in Keim's work is a mapping of discrete data points to unique pixels on the display. This typically requires distortion of the host space, something our applications cannot tolerate.

There is an extensive literature in this area (e.g., [Won97a; deO03a; Lov05a; Slo05a]). We can only scratch the surface here, and we do so using an initial categorization based on whether the technique employs a single or multiple display windows.

2.1 Multiple Windows

Visualizations that employ multiple windows may display the same geometric region in all windows, but with a different attribute (or perhaps attributes) displayed in each. The different windows might also be used as an organizational scheme for presenting how a multidimensional data space (i.e., one with $d > 3$) has been sliced down to a displayable dimension.

2.1.1 Side-by-Side Displays

In traditional side-by-side (or small multiple) displays, each display covers the same region of space, but with different attributes. (Frequently there is one attribute per map, but it is certainly possible for two or three to be present.) The user must then form a mental image of the combined maps in order to visualize all the attributes at a given location [Slo05a].

An alternative approach using a multiform bivariate matrix is used in GEOVISTA Studio for visualizing geospatially referenced objects with multiple attributes [Mac04a]. A statistic of interest (cancer mortality rates, in their example) can be correlated with pairs of possible risk factors (environmental factors, health care access, etc.). The technique employs a matrix of side-by-side displays. The map in position (i,j) , for example, is used to visualize the

extent to which risk factors i and j both lead to increased mortality rates throughout the region of study (the United States, in their case). One interesting aspect of their work is that their displays can be used to present data that is both continuously defined (e.g., presence of atmospheric emissions) and discretely defined (e.g., mortality statistics). Each map in the matrix displays only a single pair of attributes, however. The data is binned for each bivariate map, and one of a small number of colors (four, in their case) is used to visualize the resulting binned data.

2.1.2 Interacting With Multidimensional Spaces

A single graphical display of the spatial region of interest is obviously not possible when the number of independent spatial dimensions (d) is greater than 3. Typically m slices through the data must be taken (for some $m > 0$ such that $d - m \leq 3$) before a display can be generated. The Worlds within Worlds system employs a nested set of coordinate systems in one display where the placement of the i -th coordinate system inside coordinate system $(i-1)$ -st indicates how two or more of the slices was performed [Fei90a]. In addition to recording how the slices were taken, the nested coordinate systems themselves can be used as interactive data exploration tools [Mil04a].

2.1.3 Multiple Control Windows

Doleisch, et al. [Dol05] have described the use of SimVis for interactive exploration of multivariate, time-varying data. It utilizes several linked displays with different representations of attributes. This system was used to visualize the Hurricane Isabel data set we describe below. The authors in fact won first place in the visualization contest that year.

2.2 Single Display

When the entire visualization is present in a single display, we eliminate the need for the user to shift their gaze from display to display, but we pay for this by creating a generally more complex single display that presents all n attributes across the entire spatial region of interest.

2.2.1 Matrices of Attributes

An early technique similar to our Attribute Block method was used to visualize atmospheric data sets generated by the NASA Goddard Space Flight Center [Bed90a]. Hourly averages for various attributes (e.g., magnetic strength, plasma temperature, ion density) were recorded over the space of many days. A visualization was shown in which the horizontal axis was used to represent the hour, and the vertical axis was used to encode the day. Each resulting grid cell (indexed by [day, hour]) contained an array of

patches in which the color of each patch was determined by the value of a corresponding attribute. Note that the resolution of the data set was less than that of the display. Hence each grid cell from the data set could be represented by a rectangle on the display that included one colored patch for each attribute. The position of the colored patches within the grid cell were irrelevant since there was no attempt to visualize how that value may be changing across the area represented by the display cell.

While similar in appearance, the Attribute Blocks method differs in a number of respects. Most notably, it was very important to be able to decouple attribute value display resolution from the resolution at which the attribute data is stored, and to do so with minimal computational overhead. For example, we wanted to avoid the need to regenerate and resample polygonal descriptions and attribute values at each interaction during an exploratory process. We developed a very efficient approach in which we defined the geometry and all attributes once at the resolution of the original data set, and then arranged for a low-level scan conversion process to determine which attribute value to display on a pixel-by-pixel basis. Beddow's visualizations [Bed90a] can be reproduced as a special case using our method.

2.2.2 Bi- and Multivariate Choropleth Maps

Choropleth maps are a common thematic map construction in which attributes are used to determine the shading of a region of a map. Frequently the regions have geopolitical significance independent of the data; e.g., they may represent counties, states, or countries [Slo05a]. Univariate choropleth maps are probably the most common, however bivariate and even trivariate choropleth maps have been used. In the bivariate technique, two independent attributes are mapped into two distinct color scales, and the resulting colors are combined and used in the corresponding region of the map.

Approaches employing trivariate choropleth maps have also been reported. Rather than mixing three distinct color scales, texture is frequently used in lieu of one or two of the colors [Slo05a]. This technique can theoretically be extended to n attributes, but even the trivariate case can be difficult to grasp, and it must generally be used with relatively large geographic regions if textures are used to encode some attribute values.

2.2.3 Symbols and Symbol Overloading

Symbols or glyphs whose size, shape, color, topology, and/or interior structure can be systematically varied have been used to visualize (subsets of) the n attributes throughout a region of interest. Examples include stick figure icons [Pic88a], Chernoff faces, and multivariate dot

mapping techniques, to name a few. (See [Slo05a] and the extensive list of references therein.)


In some situations, we can piggyback the display of one attribute on the display of another. Vector field visualizations, for example, may be constructed from explicit arrow shapes systematically placed throughout the region, or various streamline representations constructed from numerical integrations might be produced. In either case, a common related technique is to color the resulting graphical objects according to the value of an independent scalar field also defined throughout the region. For example, the temperature at a given location could be used to color an explicit vector representation drawn at that location.

2.2.4 Perceptual Texture Elements

More recently, the basic idea of icon-based representation of multiple attributes has been extended by Healey using *perceptual texture elements* (pexels) [Hea99a]. A pexel is a polygon of fixed width whose height and color can be used to encode attribute values at a given location. Moreover, based on studies of how the human visual system perceives and analyzes texture and patterns, Healey exploits the fact that density and regularity of pexels can also be used to encode additional attributes. Three levels of regularity are identified: regular, irregular, and random. Given sufficient numbers of pexels, it is possible to perceive changes in the density of pexel placement, independent of regularity [Hea99a].

3. ATTRIBUTE BLOCK DISPLAYS

An Attribute Block display is based on interactive arrangement of all or a selected subset of the n attributes into a dynamically configurable $k_r \times k_c$ array of visual representations called an *Attribute Block* [Mil07a]. Each element of the Attribute Block array encodes a single attribute value.



(a) $k_r \times k_c = 2 \times 2$; $b_r \times b_c = 5 \times 5$ (b) Legend

Figure 1. Basic Attribute Block Structure

All attributes are defined everywhere; the Attribute Block array simply acts like a “screen door” of lenses, each allowing a single attribute value to be seen at the specific location. The dimensions of each “lens” – denoted $b_r \times b_c$ – are also dynamically adjustable. The current $k_r \times k_c$ array of $b_r \times b_c$ lenses

is then tiled across the entire region. Figure 1(a) illustrates a basic diagram of an Attribute Block configuration used to create Figure 2.

Figure 1(b) illustrates the legend displayed by the system that identifies the current $k_r \times k_c$ configuration as well as the current color ramp used to map the values for each attribute into colors. The scheme illustrated in the figure shows colors varying from the background color for minimum values of all attributes to a unique bright color for maximum values of the attributes.




Figure 2. Visualizing some parameters of a water balance Model over the upper Midwest region of the United States.

The image in Figure 2 illustrates how the four indicated parameters of a water balance model vary across the upper Midwest of the United States in the month of November. The model considers the amount of water that would be consumed, if available (potential evapotranspiration), the amount of water that can be held by the soil, the temperature in the area, and the amount of precipitation. All are monthly averages and are computed over a thirty year period [Cli02a].

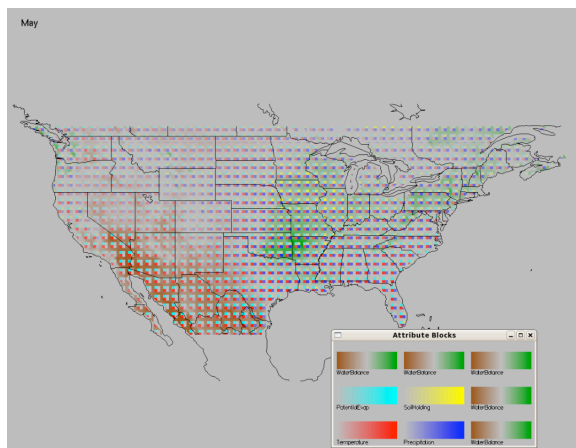


Figure 3. Visualizing the Water Balance together with four of the driving parameters of the model.

Several characteristics are immediately obvious from the figure, including the high soil moisture holding capacity in the east end of the study region, high precipitation over the western mountain region, and the high demand for water (more precisely, high potential evapotranspiration) in the south and southwest portions of the study area.

The water balance model calculates the water surplus or deficit in a region based in large part on these four parameters. Using Attribute Blocks, we can display the water balance throughout a region at the same time as we view the driving parameters. For example, Figure 3 shows the water balance (brown for deficit, green for surplus) along with these four driving parameters. Several conclusions can be immediately drawn. For example, the high soil moisture holding capacity coupled with relatively high levels of precipitation keep the water balance high in spite of relatively high temperatures and potential evapotranspiration. By contrast, the portion of the study region covering the southwest portion of the United States and northern Mexico is very dry due to poor soil moisture holding capacity and little rain along with high temperatures.

The ability to see the driving parameters interspersed with the model results combined with the ability to dynamically adjust viewing parameters such as phase shift and block size is quite effective. It allows a single large display of a region as opposed to, say, side-by-side displays which would have to be smaller and would require frequent shifting of gaze between corresponding regions of multiple displays.

3.1 Interactive Controls

Basic Attribute Block controls include the current $k_r \times k_c$ and $b_r \times b_c$ values. The display is dynamically updated as these four settings are independently adjusted. For example, pointillism-like effects are possible when $b_r \times b_c = 1 \times 1$; and the Attribute Blocks can become Attribute Rows or Attribute Columns if, respectively, k_c or k_r is set to one. Attribute Block sizes may be fixed in pixel space or model space, and dividers can be presented on the display to emphasize the boundaries between adjacent Attribute Blocks in a given display. The Attribute Blocks are a cyclic sampling of the underlying continuous fields, and one can adjust interactively the phase shift of this sampling.

In addition to these basic data-independent Attribute Block controls, several data-dependent controls are also available. These controls are all dynamically generated based on information read from the data file presented to the Attribute Block system. The legend in Figure 1(b), for example, is actually an interactive control. It is generated dynamically as the $k_r \times k_c$ settings described above are changed. Clicking

on one of the color bars of the legend produces a menu of all defined attributes, and selecting one of them results in that attribute replacing the selected one in the Attribute Block display. Various shortcuts are also included such as setting all positions (or all empty positions) to a particular attribute.

3.2 Attribute Block Display Strategies

The original motivation for the Attribute Block method was to allow n attributes to be presented in a single display by setting k_r and k_c sufficiently large so that n slots would be available. Among other things, we were interested in seeing how many attributes we could display at once in this manner and still have a display from which meaningful insight could be extracted.

One can define “meaningful insight” in different ways. For example, the ability to quickly identify regions where one attribute or another dominates is often important. Note that it may or may not be possible to infer precise quantitative measures for any given attribute, but knowing which attribute(s) is (are) most influential is meaningful. Clearly the example of Figure 2 satisfies this criterion. It is admittedly arguable whether one can precisely compute (or even define) “most influential”. It certainly depends on the nature of the data, and this is something we have been studying outside the context of this work. In any case, precise quantitative measures of parameter values are easily obtained by, for example, pointing to a specific area and having such information presented in popup windows and/or be represented with a line drawn on the legend of Figure 1(b).

It certainly comes as no surprise that a major factor affecting the understandability of the Attribute Block displays relates to the importance of color selection. The more unique attributes to be visualized at one time, the more important the selection of easily distinguishable and recognizable colors becomes. Nevertheless, that issue is not the central focus of this paper. Instead we will discuss structural aspects of Attribute Block displays that have been found to enhance the user’s ability to extract features in the data and, to a limited extent, lessen the ill effects of poor color choices.

As stated above, our original idea was to assign attributes one-by-one to positions in the Attribute Block display. What we found, however, was that strong perceptual cues could be incorporated into the display by creating certain Attribute Block arrangements in which individual attributes repeat in certain patterns.

The examples of the next several figures are visualizations of the data which modeled the

September 2003 Hurricane Isabel [IEE04a]. This data set consists of 10 scalars and 1 vector field captured at 500x500x100 resolution for 48 hourly time steps. The pressure scalar field is particularly interesting because visualizing the boundary between negative and positive pressure so clearly locates the hurricane at each time step. In these images, negative pressure is yellow; positive pressure is red; zero pressure is background (gray). For ease of comparison, all images in the paper are taken from hour 16 of this data set.




Figure 4(a): Legend for the image of Figure 4(c)

Figure 4(b): Legend for the image of Figure 4(d)




Figure 4(c): Hour 16 from the Hurricane Isabel data set




Figure 4(d): Similar to Figure 4(c), but with spaced out groups.

Figure 4 is an Attribute Block display in which two of the variables are presented for hour 16 from the 48 hour data set. As indicated in the legend of Figure 4(a), the precipitation and pressure variables are displayed in Figure 4(c) in a $k_r \times k_c = 2 \times 2$ arrangement. The image results in these two variables being grouped along diagonal lines.

By introducing some blank spots into the $k_r \times k_c$ grid, the two attributes are much more clearly separated and identified as “upper left to lower right” versus “lower left to upper right”. This is especially obvious in regular full-screen displays. (See Figure 4(b) and Figure 4(d).)

We have successfully used Attribute Blocks to visualize fairly large numbers of attributes simultaneously. For example, we see in Figure 5 an Attribute Block legend and display illustrating five attributes of the Hurricane Isabel data set.




Figure 5(a): The legend for the Attribute Block display of Figure 5(b).

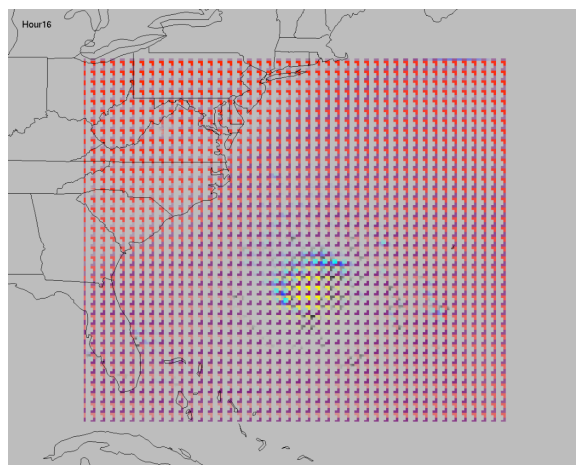


Figure 5(b): An Attribute Block display presenting the variation of five attribute values across the field of view.

As indicated in Figure 5 (a), the five attributes are (i) cloud cover, (ii) pressure, (iii) precipitation, (iv) vapor, and (v) the Graupel mixing ratio: hail, sleet, and snow pellets [IEE04a]. The latter is displayed using a cyan color ramp. This particular arrangement – and in particular, the placement of the Graupel mixing ratio in the middle of the Attribute Block pattern – seems to work well here, at least in part because the Graupel mixing ratio is heavily concentrated in some areas, but is virtually nonexistent in others.

4. DISCUSSION

As the examples in the paper indicated, we have used the Attribute Block method for two somewhat different tasks. One was to view results of a physical model along with the driving parameters of the model. The other was to enable a new way of dynamically selecting, arranging, and visualizing multiple quantities defined throughout a region. While both tasks are basically multivariate visualization, the different applications somewhat affect the strategies used. Moreover, in the former case, one of the dependent variables (water balance, in our case) was actually computed from other dependent variables. We review each application in turn and then discuss areas of further development we are pursuing to enhance the application of this technique in such application areas.

In the water balance application, the driving parameters (temperature, precipitation, soil holding capacity, and potential evapotranspiration) were known throughout a region at a certain time resolution (monthly averages, in our case). These quantities drove our computational water balance model, allowing us to approximate the water balance in the region at the same time and spatial resolution.

The original visualization of this model was developed as the key portion of a collaborative decision support system [Cli02a]. That application focused on visualizing the computed water balance as well as our level of confidence in that final measure. But to adequately address the underlying goal of the effort – collaborative policy and decision making – it became clear we needed to understand not only *where* there was a deficit or surplus of water, but also get an idea as to *why* the surplus or deficit was there. From a policy making perspective, this understanding could provide insight into appropriate remediation strategies, if any. For example, we might be able to better understand whether irrigation might be appropriate to address a periodic chronic deficit in a region, or whether instead alternative land usage decisions should be made. Coupled with the ability to incorporate climate change predictions along with their uncertainties (which the system also modeled

[Cli02a]), the tool was found to be especially helpful for policy makers [Cli02a].

In section 3.2, we discussed the possibility of using one dependent variable to implicitly group several others by making that variable dominate the Attribute Block array. As indicated in Figure 3, we found the same technique to be an effective way to cluster the display of the driving parameters of the water balance model. Among other things, it makes it easy to see the impact of soil moisture holding capacity in the upper Midwest, and the negative impact of high temperatures coupled with low precipitation and low soil moisture holding capacity in the Southwest and northern portions of Mexico. In particular, it is easy to associate (lack of) precipitation and soil holding capacity with the areas of background color that appear in the Southwest due to the grouping imposed upon the display by the Attribute Block arrangement.

The Hurricane Isabel visualization used throughout section 3.2 demonstrates the second application of the Attribute Block method. We are presented with ten scalar fields and one vector field measured throughout a study region. It certainly seems impossible to visualize at once all ten scalar quantities, hence our method is designed to support exploring these data sets by dynamically selecting a subset of the fields which can then be assigned to positions of an Attribute Block array in arbitrary ways. Tracking the value of any given scalar field across the field of view can be facilitated by an effective design of the Attribute Block layout. The ability to dynamically define the layout – both its size and the assignment of attributes to cells – has been found to enable an effective exploratory process.

While it is admittedly difficult to obtain a feel for the precise value of any given attribute at a particular location in an Attribute Block display, the strength of this method seems to be the ability to rapidly locate areas where one attribute or another is most dominant. This is certainly an important visualization task. Moreover, it is clearly straightforward to address the need for quantitative measures by supplementing this technique with a “click and probe” tool which would allow precise values to be displayed in any of a number of ways for a given attribute at a selected location.

4.1 Strategies

Effective use of this technique requires the user to exploit the interplay among the selected layout (k, xk_c), the cell sizes (b, xb_c), and the n attributes selected for presentation. The focus in the paper has been to consider Attribute Block structures that enhance the ability of the user to group given attributes to facilitate an understanding of how the attribute varies across the field of view.

Alternatively, automated tools might be developed to give advice to the user based on how individual attributes vary across the study region. For example, an attribute that does not change significantly could either be eliminated from the set of attributes to be displayed, or assigned a single cell in the layout. In general, attributes could be binned based on their variability, and the system could suggest specific layouts and initial numbers of assigned slots accordingly.

Our use of Attribute Block methods has so far been restricted to 2D domains. We are currently developing an extension of the technique that preserves its salient characteristics (e.g., decoupling of attribute display resolution from the resolution of the polygonal approximation). Included will be the ability to map Attribute Block displays onto general 3D surfaces while preserving the dynamic interactive adjustments to parameters such as the layout and cell sizes.

4.2 Tools

Additional tools outside the strict scope of Attribute Block definition and use could improve the effectiveness of Attribute Block data explorations. The line between an Attribute Block enhancement and an auxiliary tool is not always a sharp one. One such tool we are envisioning would be an intelligent color ramp generator for an attribute based on several factors including the actual distribution of the data and the other colors used in the Attribute Block layout.

4.3 User Evaluations

At this point, the system and techniques described here represent a work in progress. We have demonstrated the generality and flexibility of the framework, and we have shown how it can be used in a variety of different application domains. We have done very limited usability studies involving two domain experts and half a dozen or so knowledgeable, but non-expert evaluators. Clearly the major next step we need to take is a series of in-depth user evaluations.

A number of issues need to be studied. Included are the effects of color blending between Attribute Block cells mapped to different color ramps. The influence of the individual color ramp specifications themselves cannot be overemphasized. Several relevant issues include the fact that our eyes are not equally sensitive to all colors; colors have different meanings in different application and cultural settings; and the extent to which quantitative measures can be inferred by users from the colors is unclear. We would also like to better understand limits on how many different attributes can actually

be effectively presented at once. Strategies for designing Attribute Block layouts need to be better understood. We have noticed empirically that some configurations (involving, for example, corner sections and use of diagonally adjacent cells) makes certain patterns much more clear than do others.

5. ACKNOWLEDGMENTS

The author would like to thank Bill Kuo, Wei Wang, Cindy Bruyere, Tim Scheitlin, and Don Middleton of the U.S. National Center for Atmospheric Research (NCAR) and the U.S. National Science Foundation (NSF) for providing the Weather Research and Forecasting (WRF) Model simulation data of Hurricane Isabel.

6. REFERENCES

- [Bed90a] Beddow, J., Shape Coding of multidimensional data on a microcomputer display, *Proceedings, IEEE Visualization*, 1990, pp. 238-246.
- [Cli02a] Cliburn, D. C., Feddema, J. J., Miller, J. R. and Slocum, T. A., Design and Evaluation of a Decision Support System in a Water Balance Application, *Computers & Graphics*, Vol. 26, No. 6, December 2002, pp. 931-949.
- [deO03a] de Oliveira, M. C. F. and Levkowitz, H., From Visual Data Exploration to Visual Data Mining: A Survey, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 9, No. 3, July-September 2003, pp. 378-394.
- [Dol05] Doleisch, H., Mayer, M., Gasser, M., Priesching, P., and Hauser, H., Interactive Feature Specification for Simulation Data on Time-Varying Grids, *Proceedings of Conference on Simulation and Visualization 2005 (SimVis 2005)*, March 2005, Magdeburg, Germany, pp. 291-304.
- [Fei90a] Feiner, S. and Beshers, C., Worlds within Worlds: Metaphors for Exploring n -Dimensional Virtual Worlds, *Proceedings ACM Symposium on User Interface and Software Technology (UIST '90)*, ACM Press, 1990, pp. 76-83.
- [Hea99a] Healey, C. G. and Enns, J. T., Large Datasets at a Glance: Combining Textures and Colors in Scientific Visualization, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 5, No. 2, April-June 1999, pp. 145-167.
- [IEE04a] IEEE, *IEEE Visualization 2004 Contest*, <http://vis.computer.org/vis2004contest/>, 2004.
- [Kei04a] Keim, D. A., Panse, C., Sips, M., and NORTH, S. C., Pixel based visual data mining of geospatial data, *Computers & Graphics*, Vol. 28, 2004, pp. 327-344.
- [Lov05a] Love, A. L., Pang, A., and Kao, D. L., Visualizing Spatial Multivalued Data, *IEEE Computer Graphics and Applications*, Vol. 25, No. 3, May/June 2005, pp. 69-79.
- [Mac04a] MacEachren, A. M., Gahegan, M., Pike, W., Brewer, I., Cai, G., Lengerich, E., and Hardisty, F., Geovisualization for Knowledge Construction and Decision Support, *IEEE Computer Graphics and Applications*, Vol. 24, No. 1, January/February 2004, pp. 13-17.
- [Mac95b] MacEachren, A. M., *How Maps Work: Representation, Visualization, and Design*, The Guilford Press, New York, 1995.
- [Mil07a] Miller, J. R., Attribute Blocks: A Tool for Visualizing Multiple Continuously-Defined Attributes, *IEEE Computer Graphics & Applications*, to appear, Vol. 27, No. 3, May/June 2007.
- [Mil04a] Miller, J. R. and Gavosto, E. A., The Immersive Visualization Probe for Exploring n -Dimensional Spaces, *IEEE Computer Graphics & Applications*, Vol. 24, No. 1, January/February 2004, pp. 76-85.
- [Pic88a] Pickett, R. M. and Grinstein, G. G., Iconographic Displays for Visualizing Multidimensional Data, *Proceedings of the 1988 IEEE International Conference on Systems, Man, and Cybernetics*, August 8-12, 1988, Beijing and Shenyang, China, pp. 514-519.
- [Slo05a] Slocum, T. A., McMaster, R. B., Kessler, F. C., and Howard, H. H., *Thematic Cartography and Geographic Visualization*, second edition, Pearson Prentice Hall, Upper Saddle River, New Jersey, 2005.
- [Won97a] Wong, P. C., and Bergeron, R. D., 30 Years of Multidimensional Multivariate Visualization, In *Scientific Visualization: Overviews, Methodologies, & Techniques*, G. M. Nielson, H. Mueller, and H. Hagen (editors), IEEE Computer Society Press, Los Alamitos, CA, 1997.