Published in final edited form as:

Proc APGV. 2009 ; 2009: 61–68. doi:10.1145/1620993.1621006.

Evaluation of Glyph-based Multivariate Scalar Volume Visualization Techniques

David Feng^{*},
UNC Computer Science

Yueh Lee[†],
UNC Radiology

Lester Kwock[‡], and
UNC Radiology

Russell M. Taylor II[§]
UNC Computer Science

Abstract

We present a user study quantifying the effectiveness of Scaled Data-Driven Spheres (SDDS), a multivariate three-dimensional data set visualization technique. The user study compares SDDS, which uses separate sets of colored sphere glyphs to depict variable values, to superquadric glyphs, an alternative technique that maps all variable values to a single glyph. User study participants performed tasks designed to measure their ability to estimate values of particular variables and identify relationships among variables. Results from the study show that users were significantly more accurate and faster for both tasks under the SDDS condition.

Keywords

multivariate visualization; glyphs; user study

1 Introduction

Advances in scientific technologies have led to three-dimensional (3D), multivariate data sets becoming commonplace. Visualizing these data sets is a difficult task complicated primarily by the sheer density of information that must be projected down to an image. Regardless of the visualization technique, increasing the visibility of each scalar variable makes the resulting image more confusing.

When visualizing multivariate volume data sets, *relationship exploration* is often the primary goal. Viewers may want to verify or establish the existence of relationships among the different variables in the data set. At the same time, viewers may also wish to *estimate data values* within the space of the data. One example of such a data set comes from Magnetic Resonance (MR) spectroscopy, the area that drives our visualization study. With

©ACM, (2009).

*dfeng@cs.unc.edu

†yueh_lee@med.unc.edu

‡lester_kwock@med.unc.edu

§taylorr@cs.unc.edu

current technology, radiologists can produce approximately 20 variables in one multivariate 3D data set, of which 5-10 are of interest at one time.

In this paper, we evaluate a multivariate 3D visualization technique called Scaled Data-Driven Spheres (SDDS) that extends Bokinsky's Data-Driven spots into 3D [Bokinsky 2003]. It displays a set of spheres associated with each variable for which sphere radius corresponds to data magnitude and sphere color is used to distinguish different metabolites. By comparing the size of the glyphs, radiologists can see how two or more variables relate to each other throughout the data volume.

Radiologists studying Magnetic Resonance Spectroscopy (MRS) to better identify tumors drove our visualization design. Our radiologists colleagues use a combination of traditional tissue Magnetic Resonance Imaging (MRI) and MRS to scan brain tumors and determine tumor boundaries. Traditional MRI yields volume scalar fields that describe anatomical tissue. Radiologists combine the anatomical tissue data with spectroscopic analysis of tumors to find malignant regions in and around tumors. Radiologists often find that tumors can extend beyond what is visible in MRI, thus making accurate tumor treatment exceedingly difficult. MRS yields a volume in which each voxel contains a full metabolite spectrum. The metabolites represented in these spectra can offer a more accurate indication of the true extent of tumor growth. In our visualization system, we combine the SDDS visualization with an interactive anatomical slice plane to allow radiologists to see relationships among metabolites and anatomical features.

Performing a straightforward computation of a desired metabolite correlation field would directly answer a subset of the relationship goal: to understand the correlations among a specific set of scalar fields. To understand the relationships between more of the variables, a large number of correlation computations are necessary. The problem space grows further when other comparison metrics are used. Also, the radiologists need to see absolute metabolite concentrations in space as well relationships, which correlation fields do not address. Problems such as these require that multiple data sets be simultaneously visible so that the radiologists can gain insight for which metabolite correlations to actually test.

We show the results of a user study evaluating SDDS at the two tasks commonly required in multivariate visualizations and requested by our radiologist colleagues:

1. **Value Estimation:** What are the values of variables at particular spatial locations?
2. **Relationship Identification:** What relationships exist among variables in the data?

To see the results of these studies in the context of other techniques, we had participants perform the same tasks while viewing superquadric glyphs, an alternative multivariate 3D visualization technique. While other techniques exist for visualizing such data sets (isosurfaces, multidimensional DVR, parallel coordinates, etc.), not all of them address both visualization goals. The results of this study, presented in detail in Section 5, show that viewers are significantly more accurate and faster at both tasks when viewing SDDS visualizations.

2 MRI and MRS

Radiologists generate the MRS data set using a technique based on traditional MRI, which was derived from Nuclear Magnetic Resonance (NMR) spectroscopy. NMR was developed to probe the structure of molecules. Lauterbur and Mansfield extended these principles to provide spatially resolved information, thus creating the field of MRI [Castillo 2002]. MRI utilizes the signal from protons within the tissue of interest to produce an anatomical scalar volume data set.

Our radiologist colleagues wish to combine anatomical tissue MRI with raw spectroscopy data that describes metabolic voxel content. Analysis of spectra enables the radiologists to generate scalar volumes corresponding to the individual metabolic peaks in the spectra [Provencher 1993]. Each of these scalar volumes contains the absolute concentration of a particular metabolite. When properly understood in relation to each other, these metabolites offer a more accurate indication of the true extent of tumor growth than tissue-based imagery. This is true for two reasons. First, tumors can extend beyond the boundaries visible in anatomical MRI. By looking directly at the metabolic composition of brain tissue in addition to an anatomical signal, radiologists hope to get a better sense of the extent of a tumor. Second, distinguishing dead tissue from tissue that only appears dead is difficult after surgery, when it is crucial to know whether the tumor has been successfully removed. Directly visualizing the metabolic composition of the tissue addresses this issue by breaking the tissue signal down into individual components that radiologists can use to make this distinction.

An MR spectroscopy data set consists of approximately 20 different metabolites sampled from a human brain on a $16 \times 16 \times 9$ sample grid with voxels that are $1\text{-}2\text{ cm}^3$ in volume. Of those 20 metabolites, 5-10 are of practical interest for each particular study. In general the metabolite concentrations have a fairly low spatial frequency. This makes them amenable to isosurface rendering and volume glyphs. The radiologists also have traditional anatomical MRI of the patient's brain taken during the same sitting as the spectroscopy data. This anatomical data has a resolution of $256 \times 256 \times 112$, where each voxel is approximately $1\text{-}2\text{ mm}^3$ in volume.

3 Related Work

Several classes of techniques can be used to visualize multivariate volume data sets, including surfaces, direct volume rendering using multi-dimensional transfer functions, correlation fields, and glyphs. Abstract, non-spatial techniques are also commonly used in conjunction with spatial views. Surface-based techniques do not directly apply to this data set as they only convey data values near visible surfaces, and multiple surfaces occlude each other. Transparency reduces the occlusion while sacrificing important depth cues. We now show how the remaining classes relate to SDDS and explain which techniques were included in the user study.

3.1 Direct Volume Rendering

Direct Volume Rendering (DVR) is used to display a single scalar volume using one or more transfer functions that map data values to image voxel properties such as opacity and color. The simplest way to extend DVR to multiple variables is to separately define opacity transfer functions for each data variable and combine the resulting images using different color channels (e.g. red, green, and blue) [Cai and Sakas 1999; Rosler et al. 2006]. Images generated via this type of color mixing are difficult to interpret when more than two colors are being combined [Rheingans 1992].

Alternatively, the standard one-dimensional transfer-function can be extended to multiple dimensions. The complexity of manipulating these multi-dimensional transfer functions has resulted in active research toward user-guided semi-automatic methods [Kindlmann and Weinstein 1999; Kniss et al. 2001]. With the aid of creative user interfaces, the process of finding the right transfer conceptually becomes user-guided identification of correlations among variables.

3.2 Correlation Fields

Many techniques exist for highlighting potentially meaningful correlations in multivariate data to users. Nattkemper reviewed the application of several of these techniques to biomedicine, some of which we describe below [2004]. Principal component analysis and other dimensionality reduction approaches attempt to project higher dimensional data spaces down to two or three orthogonal dimensions that represent the greatest variation in the data. Broersen and van Liere apply PCA to raw spectroscopy data to find images that represent the greatest variation in both feature space and spectrum space [2005]. They subsequently auto-generate opacity transfer functions for the computed eigenvectors and display the images using standard DVR.

Our collaborators require an exploratory view to produce hypotheses about which relationships matter. Crouzil et al. describe an interface that uses gradient alignment of different scalar fields to display multiple correlations [1996]. Multi-field graphs provide an exploratory interface to the large number of potential correlations and use DVR to display a volume of interest in those correlations [Sauber et al. 2006]. Woodring and Shen have designed a system in which users can combine an arbitrary number of scalar fields using various set operations (e.g. AND, OR, XOR, etc.) to generate an interactive expression tree [2006]. These correlation detection routines turn the multivariate visualization problem into a guided search through all possible relationships of which there are $O(n!)$ (where n is the number of fields)..

These techniques result in a single compressed data field. Our radiologist colleague's must see the raw values of individual fields in addition to the relationships among them, so we focus on a visualization techniques that can display multiple fields at once.

3.3 Sparse Glyphs

In sparse glyph visualization techniques, data values are represented via the properties of geometrical shapes (glyphs). The glyphs must be large enough that multiple data values can be represented at once using multiple parameterizations that map properties like shape, size, color, and opacity to variable data values.

The use of sparse volume glyphs for visualization has been actively studied in tensor field visualization. Kindlmann et al. describe the use of superquadrics and other shapes for glyph-based tensor visualization [2006]. They vary the shape and orientation of glyphs to indicate the magnitude and direction of flow at glyph locations. Their work is primarily based in 2D tensor field visualization in which glyph shape varies according to tensor components. Forsell et al. used 3D glyphs on multivariate data in which they vary properties like specularity and concavity of a 2D surface representing a 2D scalar field [2005]. These surfaces do not easily extend into 3D.

Ebert et al. have studied glyph usage for multi-dimensional data visualization, primarily by discussing the different ways of varying shape to convey different scalar values [2000]. They propose varying color, size, shape, and opacity along separate scalar components. Such an encoding is problematic for the MR spectroscopy data set because using two different encodings for two semantically similar metabolite fields makes relative magnitude estimation difficult. Some variables will be harder to interpret than others. To address these issues, we use only color to differentiate metabolite volumes. The perceptual differences among colors are smaller than those among shape encodings, as is confirmed in the user study presented in Section 5.

4 Scaled Data-Driven Spheres

SDDS is a 3D extension of Bokinsky's 2D Data-Driven Spots, a 2D multivariate visualization technique [Bokinsky 2003]. In her work, Bokinsky displays multiple scalar fields using color-encoded Gaussian splats placed on a jittered sample grid and shows that multiple layers of differently-colored spots were as effective for the display of the shape of overlapping 2D scalar fields as direct display of the computed intersection. SDDS extends this to 3D while avoiding opacity color mixing. BrainExplorer, developed by Lau et al., uses a similar sphere-based glyph technique to visualize gene expression in mouse brains [Lau et al. 2008]. This technique was developed concurrently to SDDS and maps glyph size to expression level and glyph color to anatomical annotation, gene type, or gene expression level redundantly. SDDS is similar to BrainExplorer when glyph size is mapped to expression level and glyph color separates different genes. The results of the user study described in Section 5 should apply to the BrainExplorer work as well SDDS.

The SDDS technique distributes shaded 3D glyphs throughout the brain volume, as shown in Figure 3. We use spheres for the glyphs, as these simple glyphs are easy to interpret at a wide range of scales. Spheres placed along a regular grid exhibit strong aliasing effects, so we resample the metabolite concentration fields on a jittered version of the original sample grid. Each scalar volume uses a separately generated jittered grid. The sphere radius is determined as follows:

$$r_{max} = k \frac{s}{2n} \quad (1)$$

$$r = r_{max} \frac{v - v_{min}}{v_{max} - v_{min}} \quad (2)$$

where r_{max} is maximum glyph radius, r is the glyph radius for a particular data value v in the range $[v_{min}, v_{max}]$, s is the sample spacing, n is the number of scalar volumes visible and k is a user-adjustable parameter. When $k = 1$, the glyphs for all scalar volumes at a single sample point can fit within a voxel without over-lapping. When sphere glyphs get too large, they begin to occlude each other, hiding portions of the volume.

We display multiple metabolite concentration fields by labeling multiple sets of spheres with unique color values. Because the human visual system perceives color preattentively, viewers can easily distinguish the nominally colored variables. In her 2D Data-Driven Spots visualization, Bokinsky found that viewers could visually attend to a single field in the presence of at least 8 other fields [Bokinsky 2003]. The human visual system struggles when differentiating more than 12 grouped color values [Ware 2000; Healey 1996]. Barring the background color, 11 thus serves as the theoretical upper limit of the number of simultaneously renderable metabolite concentration fields. Selecting isoluminant colors for the glyphs would prevent one variable from appearing easier to locate than another, but results from our user study, described in Section 5, show that the measured difference for these tasks is not statistically significant.

The inevitable increase in visual occlusion resulting from adding a new field also limits the number of renderable fields. Depending on the desired glyph size, sample density, and data sparsity, rendering more than five to six scalar fields simultaneously begins to cause over-occlusion. Figure 3 contains an SDDS visualization in front of an anatomical slice plane,

showing the correlation between choline concentration (yellow spheres) and lesion extent (gray anatomy).

4.1 Anatomical Slice Planes and Stereo Viewing

Radiologists must see the spectroscopy data in the same space as anatomy data to evaluate tumor activity in and around lesions visible in anatomical scans. We therefore combine SDDS with an interactive gray-scale anatomical slice plane.

Stereo helps viewers to disambiguate depth in the renderings and thereby have a fuller understand the shape of the different scalar fields. This is primarily needed due to the large amount of data on display at any given point in time. Stereo is often used in visualizations of amorphous, organic structures, where shape and relative position of structures is important. Molecular visualization is the most common of these types of visualizations, for which researchers frequently include cross-eyed or anaglyphic stereo pairs for print publications. Shutter stereo goggles, which synchronize with graphics hardware to display different images to each eye, have become well-supported, inexpensive hardware components that are easy to integrate into any 3D application.

5 Glyph User Study

We designed a study to evaluate SDDS on how well it addresses the MR spectroscopy visualization goals, renumerated below:

1. **Estimate values** of multiple variables in local data regions.
2. **Identify relationships** between data variables.

The study therefore consisted of two separate tasks performed by participants that explored both goals.

While absolute performance measures are useful, to put the results of this study in context with other similar visualization techniques, we therefore also evaluate superquadric glyphs. We compare SDDS only to superquadric glyphs because we are unaware of any other visualization techniques that can simultaneously display the raw values of four or more volume fields. DVR with multi-dimensional transfer functions (Section 3.1) and guided exploration techniques (Section 3.2) reveal relationships, but they are generally used to visualize a single field representing a computed relationship of interest. Surface-based techniques described in Section 3 can be used to display multiple fields, but surfaces do not convey raw data values in regions away from the surface. 3D glyph-based techniques are the strongest candidates for this type of visualization, and we use superquadric glyphs to represent the class of techniques that represent multiple variables in a single, more complex glyph.

We generated all of the visualizations in the study using randomly generated four variable data sets. We chose this number of variables for two reasons. First, our radiologist colleagues are currently interested in four metabolites (choline, creatine, NAA, and lipids). Second, toroidal superquadric glyphs have four natural variable properties: thickness, overall roundness, cross-section roundness, and color. SDDS can potentially accommodate more than four variables, but additional properties for the superquadric glyphs (e.g., size) tend to interfere with the other properties. For the spheres, we assigned the variables to four sphere colors: red, yellow, green, and cyan. These match the basic color opponency channels in the human visual system.

The range of glyphs for each variable property of the superquadric glyphs was chosen to maximize dynamic range and intuitiveness. Toroidal thickness ranged from thin but visible

to thick but open. The color variable used a truncated black-body radiation color map with black removed so as to not conflict with the dark background. This resulted in a red-yellow-white color map.

There is a clear trade-off between dynamic range and intuitiveness in the roundness variables. We decided that they should vary from completely square to completely round, even though both variables can continue beyond completely round shapes to produce concave shapes. We did not include this range of shapes because the transition from round to concave is not as intuitive as the transition from square to round. A shape map including both sets of transitions contains a perceptual discontinuity that does not correspond to features in the data.

We sampled the data using four SDDS spheres for every superquadric glyph, where the spheres each had diameters of approximately half the width of the superquadric glyphs. This ensured that both glyph techniques sampled the data at the same density. As a result, the SDDS spheres were four times as dense as the superquadrics.

The randomly generated collection of data sets were designed to have similar properties to MR spectroscopy data. The MR spectroscopy metabolite fields exhibit slow spatial variation, so each data set was a composite of several Gaussian fields with randomized centers, orientations, and variances in each dimension. The four data sets for each trial were randomly selected from a set of 50 such fields. The randomized order of data sets was different for each participant. After randomly selecting an initial visualization type, visualizations alternated between SDDS and superquadrics.

User study participants were asked to view the glyph visualizations while wearing stereo goggles. They were also given the ability to toggle movement and rotation of the visualization camera along a fixed path. We did not allow users to directly manipulate the camera (for example, using the standard mouse/trackball interaction style) because a participant's previous experience with the mouse and user interface could potentially confound results.

5.1 Task 1: Value Estimation

For the first task, participants viewed alternating visualizations using SDDS and superquadric glyphs and were asked to estimate the value of two of the four visible variables at a particular region in space. Participants were allowed to enter discrete multiples of .1 ranging from 0 to 1. The region of interest was labeled using a white wireframe cube. Because this is an inherently 3D task, the visualizations also contained a 3D legend for the two variables of interest that rotated along with the glyphs, as shown in Figure 4. In the superquadric legend, all variables except for the variable of interest were set at their middle value (0.5). For this task, we measured response accuracy and response time.

5.2 Task 2: Correlation Identification

The second task consisted of participants identifying the two variables present in the data that contained a strong positive correlation. The data for three of the variables were randomly selected from a data set generated as described above. One of those three variables was subsequently selected at random to be uniformly scaled down. This then became the fourth data set, thereby ensuring that any accidental positive correlations of the other variables would not be as strong as the perfect correlation. As with Task 1, a 3D legend was presented to the users, however this legend displayed all four variables at once, as shown in Figure 5. For this task, we tracked the participant's two responses and response time.

5.3 Equipment and Materials

We ran the user study on a computer containing a NVIDIA Quadro FX5600 GPU with 3-pin stereo output. Participants wore CrystalEyes goggles throughout the experiment synchronized with a 21" flat screen CRT refreshing at 60 Hz per eye. All interaction with the visualization was through a 40-key X-Keys programmable keypad customized for this study. We chose this input device to avoid potential error caused by varying levels of computer experience in the participants. The study was performed in a darkened room with a small desk lamp illuminating the keypad.

5.4 Procedure

We ran the study on 17 participants (14 male, 3 female). All participants were first led to a private room where they read a short document explaining the data and visualization techniques. They then read about the first task (value estimation) and began an 8 trial training session with no data recorded, mimicking the upcoming recorded session. Participants could ask clarifying questions throughout the introduction and training session. Once the training session was complete, the tester left the room and participants began a 60 trial recorded session on the first task.

Participants could take a short break following their completion of the first task, after which they read a short introduction to the second task (correlation identification). A second 8 trial training session then began with the tester present, followed by 40 timed trials with the tester absent. We used fewer trials for this task both to avoid fatigue in participants and because a pilot study indicated that 40 trials would be sufficient to achieve statistical significance. Once participants finished the second session, they filled out a short follow-up questionnaire asking them for visualization preference, ratings for each task, and general comments.

6 User Study Results

We now present the accuracy and response time for both user study tasks. Response error for Task 1 and response times for both tasks were analyzed using generalized linear models with normal distributions. Because Task 2 error is a binary value (correct vs. incorrect), we used a generalized linear model with a binary distribution. All error bars in the figures correspond to standard error. Additionally, participants did not exhibit a significant learning effect in either the value estimation task or the correlation identification task. The initial training sessions for each task successfully help participants past the initial learning curve.

6.1 Task 1: Value Estimation

Participants viewing the SDDS condition responded with an mean error of .0839 units, or 8.39% of the value range. Because participants could only enter discrete multiples of .1, perfect accuracy would involve an expected mean error of 0.05, or 5% of the value range. The mean response time for SDDS was 27.7 seconds.

Response error was on average lower for the SDDS condition. The mean value estimation error for superquadrics was .1325 units (13.35% of the value range). Participants were 37% more accurate with the SDDS condition. The probability that these two means represent the same distribution (the p value) is less than .0001. Figure 6 presents this data along with standard error bars.

Response time was also lower for the SDDS condition. The mean response time for superquadrics was 34.5 seconds, a 20% mean improvement from SDDS to superquadrics. This is the measured time it took a participant to estimate the values of two variables, so we

infer than the mean response time for estimating a single value is 13.4 seconds for SDDS and 17.3 seconds for superquadric glyphs.

Whereas the error across the different SDDS colors was not statistically different, the value estimation error for the superquadric glyphs was weighted heavily to the cross-section roundness variable and to a lesser extent the overall roundness variable. The mean error for the SDDS colors confirms that these four colors, despite having different luminance, do not significantly affect viewer performance. The mean error per variable is shown in Figure 8.

6.2 Task 2: Correlation Identification

For the SDDS condition, participants incorrectly identified the correlated pair of variables on 23.8% of the trials with a mean response time of 23.1 seconds. Participants viewing the superquadric condition incorrectly identified the correlated pair 79.1% of the time with a mean response time of 42.8 seconds. Comparatively, participants were 70% more accurate and 47% faster under the SDDS condition, both with $p < .0001$. This comparison is shown in Figure 9. The error bars on Figure 9 are asymmetric because binary distribution analysis is done in log space.

Participants viewing the superquadric glyph condition correctly identified the correlated pair only slightly more often than chance. For four variables, the probability of randomly selecting the correct pair is 16.7%; participants correctly identified the correlated pair only 20.8% the time. Figure 11 shows the mean frequency with which participants selected a variable that did not belong to the correlated pair. The four SDDS color variables were chosen incorrectly at a uniform rate, whereas the error for the superquadric properties is unevenly distributed. Participants tended to select color for the superquadric glyphs far more often than they should have. Responses from the follow-up questionnaire confirm the fact that participants often began looking at color first because it was the easiest property to see. This imbalance in shape properties has a significant effect on what variable relationships viewers perceive.

One concern with SDDS is that the sphere colors have different luminance and thus are perceived differently. For the correlation identification task, the differences between mean identification errors for the four SDDS colors do not appear to vary with color luminance, nor do they appear to vary significantly.

6.3 Follow-up Questionnaire

After completing Task 2, participants filled out a short questionnaire in which we asked them to evaluate the two visualization techniques they saw. All participants but two preferred the SDDS visualizations over the superquadric glyph visualizations. Reasons for these responses included the simplicity of the color map and value range of the spheres and the complexity and variability of the superquadric properties. One participant preferred the superquadric glyphs for Task 1 because the glyphs were less densely packed, but found SDDS simpler for Task 2. The data for the two participants did not support their stated preferences, as both performed more accurately and faster when viewing SDDS visualizations.

We also asked participants to rank the four superquadric variable properties in order of ease of interpretation. The rank for each variable are as follows: color ranked 1, thickness ranked 2.23, overall roundness ranked 2.77, and cross-section roundness ranked 4. Participants were in unanimous agreement about the rankings of color as easiest and cross-section roundness as most difficult. Several participants made the comment that the color and thickness variables “stood out” to them, whereas the roundness variables did not. Most participants

commented on the difficulty of interpreting the cross-section roundness variable. These perceived preferences support the error measurements taken in Task 1 and Task 2.

Participants also made several recurring general comments about the different glyphs. Many participants required repeated explanations to understand the cross-section variation, and many also said that this property was difficult to understand because of insufficient shape variability. This is reflected by the individual variable errors shown in Figures 8 and 11. Increasing this channel's dynamic range by including concave shapes could ameliorate this discrepancy.

7 Conclusion and Future Work

We presented a user study evaluating the effectiveness of SDDS for multivariate 3D data sets similar to those used by radiologists studying MR spectroscopy. Effectiveness has been defined as how well viewers could perform two tasks designed to address the original visualization goals: value estimation and relationship identification. Additionally, we have evaluated superquadric glyphs, alternative spatial multivariate 3D visualization technique, for context and comparison. While the error rates were nontrivial, they are not unreasonable for an exploratory visualization system. Compared to the superquadric condition, participants were significantly more accurate and faster at identifying positive correlations and estimating data values with SDDS visualizations. The contrast between the two visualization techniques was particularly dramatic for the correlation identification task, with participants responding nearly twice as quickly and more than twice as accurately.

We do not claim to have tested the ideal set of shape-varying glyph properties or the perfect dynamic range for those properties, nor do we claim that SDDS is always better for these tasks. However, the results do highlight an important effect of choosing variable channels that are not perceptually equivalent. The most common explanation given by participants for why they preferred SDDS to superquadrics was that the superquadric roundness properties were much more difficult to interpret than thickness or color, which was the property that participants understood the most accurately. SDDS uses color in a more perceptually uniform nominal color encoding for different variables. Because humans perceive color preattentively, distinguishing between two colors is generally easier than distinguishing between color variation and shape variation.

From a visualization design perspective, this difference between variable channels results in an emphasis of one variable over another. This is not necessarily a bad decision. The freedom to manipulate the dynamic range of particular variables that results from using perceptually different channels is useful if one variable is more important than another for the viewer. That said, a study of the parameter space of shape-varying glyphs is necessary to customize dynamic range accurately. This requirement may prohibit the use of significantly different variable channels in a visualization like the superquadrics described in this work.

Glyph-based techniques like SDDS apply most directly to data sets with slow spatial variation, as is the case with the MR spectroscopy data set. Applying SDDS to high-frequency data sets will reveal low frequency trends in properly filtered data. SDDS will also work well when visualizing sub-regions of high frequency data.

The SDDS visualization enables viewers to explore and analyze relationships in multivariate volume scalar fields. To our knowledge, SDDS is the only multivariate scalar volume visualization technique that can potentially scale to 11 simultaneous display channels. Multivariate scalar visualization problems also exist in confocal microscopy, where researchers capture volume data labeled with multiple fluorophores. SDDS may also apply in cell chemistry simulation, where complex cell reactions are observed over time. As we

continue to improve the SDDS technique, we hope to enable such researchers outside of MR spectroscopy to better understand the relationships and values in their complex data sets.

References

- Bokinsky, A. PhD thesis. UNC; Chapel Hill: 2003. Multivariate Data Visualization with Data-Driven Spots.
- Broersen, A.; van Liere, R. EuroVis. Eurographics Association; 2005. Transfer functions for imaging spectroscopy data using principal component analysis; p. 117-123.
- Cai W, Sakas G. Data intermixing and multi-volume rendering. *Computer Graphics Forum*. 1999; 18:359–368.
- Castillo, M. *Neuroradiology*. Lippincott Williams & Jenkins; 2002.
- Crouzil, A.; Massip-Pailhes, L.; Castan, S. A new correlation criterion based on gradient fields similarity. *International Conference on Pattern Recognition*; 1996. p. 632-636.
- Ebert DS, Rohrer RM, Shaw CD, Panda P, Kukla JM, Roberts DA. Procedural shape generation for multi-dimensional data visualization. *Computers & Graphics*. 2000; 24:375–384.
- Forsell, C.; Seipel, S.; Lind, M. Simple 3D glyphs for spatial multivariate data. *IEEE Symposium on Information Visualization*, 2005; 2005. p. 119-124.
- Healey, CH. Choosing effective colours for data visualization. *Proceedings of IEEE Visualization '96*; 1996. p. 263-270.
- Kindlmann, G.; Weinstein, D. Hue-balls and littersors for direct volume rendering of diffusion tensor fields. *Proceedings of IEEE Visualization '99*; Los Alamitos, CA, USA: IEEE Computer Society Press; 1999. p. 183-189.
- Kindlmann G, Westin C-F. Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization and Computer Graphics*. 2006; 12(5):1329–1335. [PubMed: 17080869]
- Kniss, J.; Kindlmann, G.; Hansen, C. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. *Proceedings of IEEE Visualization '01*; IEEE Computer Society; 2001. p. 255-262.
- Lau C, Ng L, Thompson C, Pathak S, Kuan L, Jones A, Hawrylycz M. Exploration and visualization of gene expression with neuroanatomy in the adult mouse brain. *BMC Bioinformatics*. 2008; 9:153. [PubMed: 18366675]
- Nattkemper TW. Multivariate image analysis in biomedicine. *J. of Biomedical Informatics*. 2004; 37(5):380–391.
- Provencher S. Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magn Reson Med*. Dec.1993 30:672–679. [PubMed: 8139448]
- Rheingans, P. Color, change, and control for quantitative data display. *Proceedings of IEEE Visualization '92*; Los Alamitos, CA, USA: IEEE Computer Society; 1992. p. 252-259.
- Rösler, F.; Tejada, E.; Fangmeier, T.; Ertl, T.; Knauff, M. GPU-based multi-volume rendering for the visualization of functional brain images. *Proceedings of SimVis 2006*; 2006. p. 305-318.
- Sauber N, Theisel H, Seidel H-P. Multifield-graphs: An approach to visualizing correlations in multifield scalar data. *IEEE Transactions on Visualization and Computer Graphics*. 2006; 12(5): 917–924. [PubMed: 17080817]
- Ware, C. *Information visualization: perception for design*. Morgan Kaufmann Publishers Inc.; San Francisco, CA, USA: 2000.
- Woodring J, Shen H-W. Multi-variate, time varying, and comparative visualization with contextual cues. *IEEE Transactions on Visualization and Computer Graphics*. 2006; 12(5):909–916. [PubMed: 17080816]

Figure 1.

Example visualizations shown to participants asked to perform two tasks in a user study comparing two multivariate 3D visualization techniques. Left: a scaled data-driven spheres visualization depicting a 4 variable synthetic data set for which participants estimated the value of variables at a location annotated by a wireframe cube. Right: a superquadric glyph visualization of a similar data set for which participants identified the two positively correlated variables. See Sections 4 and 5 for details on the visualization techniques and study design.

Figure 2.
Our radiologist colleagues currently visualize the MRS data set by overlaying raw metabolite spectra on top of anatomical MRI slices. This visualization is difficult to understand, even for specially trained radiologists.

Figure 3. SDDS visualization applied to a data set containing a visible lesion, shown in gray in the background anatomical data. In this example, the yellow and orange spheres correspond to choline concentration and creatine concentration, respectively. Using SDDS, it is apparent that the yellow and orange spheres have an inverse relationship outside the lesion, but both decrease in the lesion. The relationships is significantly more apparent in 3D with stereo and motion cues, as described in Section. 4.1

(a) Superquadric Shapes

(b) Task 1 Example

Figure 4.

a) Example superquadric shapes with four variable properties. b) An example visualization of the type participants saw during the value estimation task. The participant first estimated the value in the cube for variable A, as designated by the labeled 3D legend, then variable B.

Figure 5.

Example visualizations of the type participants saw during the correlation identification task. The participant estimated which two variables, demonstrated in the 3D legend above the data, were positively correlated with each other. The same data is used in both visualizations, where variables B and C have a strong positive correlation. When viewed in stereo and in motion, the images are sparser and clearer.

Figure 6.

Average value estimation error for the first user study task. Users were 37% more accurate with the SDDS visualization as compared to the superquadric glyph visualization. Expected error due to discretization of possible responses is .05.

Task 1: Value Estimation Response Time (p=0.0042)

Figure 7.

Average response time for the first user study task. Users were 20% faster with the SDDS visualization as compared to the superquadric glyph visualization.

Figure 8.

Average error all glyph properties. Cross-section roundness and to a lesser extent overall roundness accounted for most of the error for the superquadric glyphs.

Figure 9. Average correlation identification error for the second user study task. Users correctly identified the correlated pair 70% more often with the SDDS visualization as compared to the superquadric glyph visualization.

Figure 10.

Average response time for the second user study task. Users were 47% faster with the SDDS visualization as compared to the superquadric glyph visualization.

Figure 11.

Average number of incorrect identifications per glyph property for the second user study task. The sphere glyph properties all exhibit uniform error distributions. The color property of the superquadric glyphs was chosen incorrectly noticeably more often than the other shape properties.