Published in final edited form as:
Vis Data Anal. 2010 January 18; 7530: . doi:10.1117/12.839818.

Linked Exploratory Visualizations for Uncertain MR Spectroscopy Data

David Feng, **Lester Kwock**, **Yueh Lee**, and **Russell M. Taylor II**
University of North Carolina at Chapel Hill

Abstract

We present a system for visualizing magnetic resonance spectroscopy (MRS) data sets. Using MRS, radiologists generate multiple 3D scalar fields of metabolite concentrations within the brain and compare them to anatomical magnetic resonance imaging. By understanding the relationship between metabolic makeup and anatomical structure, radiologists hope to better diagnose and treat tumors and lesions. Our system consists of three linked visualizations: a spatial glyph-based technique we call Scaled Data-Driven Spheres, a parallel coordinates visualization augmented to incorporate uncertainty in the data, and a slice plane for accurate data value extraction. The parallel coordinates visualization uses specialized brush interactions designed to help users identify nontrivial linear relationships between scalar fields. We describe two novel contributions to parallel coordinates visualizations: linear function brushing and new axis construction. Users have discovered significant relationships among metabolites and anatomy by linking interactions between the three visualizations.

Keywords

medical visualization; multivariate data; parallel coordinates; glyphs; linked views; uncertainty visualization

1. Introduction

We present a system for visualizing magnetic resonance spectroscopy (MRS) data in the brain. Radiologists use the concentrations of multiple metabolites and their interrelationships to understand brain physiology. This understanding is important for identifying abnormalities such as tumors that are not always clear using anatomical magnetic resonance imaging (MRI). For example, tumors can extend beyond the boundaries visible using traditional MRI techniques. Also, radiation necrosis (dead tissue resulting from radiation treatment) is difficult to distinguish from recurring tumors after surgery. MRS reveals the chemical composition of tissue, which enables radiologists to more easily distinguish between malignant and benign tissue.

Our collaborators have two primary visualization goals:

1. **Relationship Exploration:** Radiologists use both absolute metabolite concentrations and relationships among metabolites as indicators of disease

processes. Our colleagues seek to confirm known indicators and identify new metabolite relationships for use in diagnosis and treatment.

2. **Surgical Planning:** Once radiologists identify a tumor, clinicians and surgeons need to extract data values at precise locations to accurately plan procedures. This goal requires viewers to have positional awareness within the data space and direct access to raw data values.

The MRS data is composed of spatially registered anatomical brain MRI and spectroscopy images. Anatomical MRI resolves different tissues such as gray matter and white matter as different scalar intensities. Voxels in anatomical MRI are ~ $1mm^3$ in volume. To generate spectroscopy data, a statistical analysis tool identifies peaks in metabolic spectra with a measured degree of uncertainty. Each voxel has over 20 metabolite concentrations represented by a normal distribution with mean and standard deviation. Spectra are captured at ~ $1cm^3$ resolution. The coarse resolution allows us to use spatial visualization techniques that may be confusing for higher resolution data.

Our visualization system has three components. First, we show users a 3D spatial visualization called Scaled Data-Driven Spheres (SDDS). Using SDDS, a sphere scaled by concentration and colored uniquely for each metabolite represents each data sample. We randomly offset the spheres from their sample position and scale them such that spheres from all variables are distinguishable. An interactive anatomical slice plane gives context to the spatial visualization. This visualization helps our colleagues get a global sense of spatial relationships among metabolites and anatomical structure.

The second visualization is a parallel coordinates (PC) visualization designed to help radiologists confirm and explore hypotheses generated via the SDDS visualization. A PC plot is an abstract data representation that plots data samples (for MRS, voxels) as lines traveling through multiple parallel variable axes. PC is most useful for identifying patterns, such as clusters of values or relationships. We have augmented traditional PC to convey data uncertainty by representing curves directly as Gaussian distributions and implemented two novel interaction techniques:

- **Linear Function Brushing**: Custom brush interactions identify linear relationships between pairs of variables.
- **New Axis Construction**: User-identified relationships can be mapped into new PC axes, which can serve as the basis for more complex relationship discovery.

The strengths of the SDDS and PC visualizations complement each other. The density of information in an SDDS visualization can result in clutter and occlusion problems, but the PC visualization compensates by discarding spatial information. When a viewer finds patterns in the PC visualization, they can refer to the SDDS visualization to understand the positions of interesting voxels and their relation to anatomy. We link interactions in both visualizations so that data points selected in one visualization are reflected in another.

The third visualization component satisfies the surgical planning visualization goal. Once users have identified a meaningful relationship or set of voxels, our colleagues need an interface for precisely extracting MRS data values in anatomical context. An interactive pseudocolored slice plane with isovalue contours shows values of a particular metabolite overlaid on a grayscale coloring of anatomy. While simple, this visualization is a critical last step before information discovered in the exploratory visualizations can be used in practice.

2. MRI and MR Spectroscopy

Our radiologist colleagues 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.[1] MRI utilizes the signal from protons within tissues of interest to produce an anatomical scalar 3D data set. The anatomical images shown in this work were generated via T1 (or spin-lattice) relaxation imaging. T1 represents an interaction between a proton and local tissue as characterized during the recovery of a disturbed magnetic field to equilibrium. Similarly, MRS uses principles of NMR spectroscopy to probe the underlying metabolic spectrum of tissue.[2]

The raw spectroscopy data consists of per-voxel metabolite spectra in which the heights of different spectral peaks correspond to different metabolite concentrations. Each spectrum is normalized on a per-voxel basis, which means that extensive expertise is necessary to properly understand peak correspondence and metabolite interactions between voxels. An offline processing system called LC Model computes absolute metabolite concentration data sets based on the normalized samples.[3] For each resolvable metabolite peak, LC Model generates a scalar volume data set. Each absolute voxel concentration measurement is made via comparison to pure metabolite basis sets, so LC Model also computes an appropriate standard deviation with respect to established normal values. Depending on the imaging sequence, LC Model generates at least 20 different metabolite concentration scalar volumes. The data sets commonly associated with brain spectra, including choline, creatine, inositol, glutamine, and N-acetylaspartate, are selected for further processing and visualization.

The current state of the art in medicine for handling MRS does not display multiple metabolite concentration fields at once. Our radiologist colleagues overlay metabolite spectra over each voxel of a slice through the anatomical data, as shown in Figure 1. Maudsley et al. use a spectrum pseudocoloring to encode a single metabolite's concentration on top of a gray scale anatomical image.[4] Chang et al. use the computed ratio of choline to creatine to diagnose gliomas, a particular type of brain tumor;[5] they superimpose a grayscale pseudocolored computed field over a grayscale anatomical slice plane. Our radiologist colleagues require a technique for displaying multiple metabolite fields at once while supporting relationship identification rather than displaying a single metabolite relationship. Radiologists capture both brain MRI and MRS images using a Siemens Allegra 1.5T MRI scanner. Each acquisition session lasts approximately six minutes. The metabolite spectra are sampled at a voxel size of $\sim 1cm^3$, with resolution on the order of 20×20×10. Because both anatomical MRI and MRS data sets are captured simultaneously for a single patient, no additional registration is necessary.

3. Related Work

SDDS is a 3D multivariate scalar visualization technique that uses sparse 3D glyphs to represent data values. Research in visualizing such data has approached the problem via alternative techniques such as direct volume rendering (DVR), computed correlation fields, isosurfaces with glyphs, and sparse 3D glyphs. The PC visualization pulls from a long history of work most recently started by Inselberg.[6]

3.1 Alternatives to Sparse 3D Glyphs

Direct Volume Rendering (DVR) displays a scalar field using one or more functions that map data values to voxel properties such as opacity and color. DVR can extend to multivariate data when users define multiple transfer functions and combine the resulting

images using different color channels.[7] However, the resulting color mixing is difficult to interpret when more than two colors are combined.[8] Multidimensional transfer functions address this by creating a higher-dimensional transfer function of multiple variables.[9] With creative user interfaces, the process of finding the right transfer function is analogous to user-guided identification of variable correlations.

There are many techniques that highlight potentially meaningful correlations in multivariate data. Nattkemper reviewed several of these techniques as applied to biomedicine, some of which we describe below.[10] Broersen and van Liere apply principal component analysis (PCA) to raw spectroscopy data to find images that represent the greatest variation in feature space and spectrum space.[11] They subsequently auto-generate opacity transfer functions using the PCA eigenvectors for use with standard DVR. While often useful, greatest variance is only one of many useful descriptors and is not the relationship of interest to our radiologist colleagues.

Our collaborators require an exploratory tool to help them to form hypotheses about which metabolite relationships matter. Crouzil et al. describe an interface that uses gradient alignment of scalar fields to display strong correlations.[12] Multifield graphs let users explore the large space of possible correlations and use DVR to display correlations of interest.[13] Woodring and Shen describe a system that lets users combine scalar fields using set operations (e.g. AND, OR, XOR, etc.) to generate an expression tree.[14] Our work complements these by giving users a sense of the raw data values lost after relationship computation. The SDDS visualization enables the radiologists to explore relationships among 5-10 fields, whereas computed correlations show radiologists specific, complex combinations of fields. We incorporate this ability by enabling users to construct new axes in the PC visualization as they search for new relationships.

Surface representations are most useful in applications where object segmentation and shape visualization are the goals. To extend surface-based techniques to multivariate data, the simplest option is to create separate isosurfaces for each variable and render them simultaneously. Because opaque surfaces occlude each other, some form of surface transparency is necessary. However, Interrante et al. showed that the human visual system does not interpret depth accurately behind transparent surfaces.[15] Partially transparent textures, often based on principal surface curvature directions, help address this issue.[16] While such surface-based techniques can be useful for visualizing several variables at once, they do not address the visualization goals for our application. Viewers cannot reliably estimate data values using surface visualizations except in areas near the surface. Our visualization goals require a technique that both preserves data value visibility and shows variable relationships.

3.2 Sparse 3D Glyphs

In sparse glyph visualization techniques, variable values are represented via separate properties of geometrical glyphs (shape, size, color, opacity, etc.). The glyphs must be large enough that variable channels are distinguishable. Kindlmann et al. describe superquadrics and other shapes for glyph-based tensor visualization.[17] They vary the shape and orientation of glyphs to indicate the magnitude and direction of flow at glyph locations.

Ebert et al. studied glyph usage for multi-dimensional data visualization, primarily by discussing the different ways of varying shape to convey different scalar values.[18] They propose varying color, size, shape, and opacity along separate scalar components. Such an encoding is problematic for the MRS data set because using different encodings for semantically similar metabolite fields makes relative magnitude estimation difficult. Some variables will be harder to interpret than others. Also, varying shape for one variable and

varying size for another does not convey the impression that all scalar fields have the same modality. To address these issues, we use only color to differentiate metabolite volumes.

This paper presents a 3D glyph technique that extends Bokinsky's Data-Driven Spots into 3D.[19] Bokinsky displays multiple scalar fields using color-encoded Gaussian splats placed on a jittered sample grid. She 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. BrainExplorer, developed by Lau et al., uses a similar sphere-based glyph technique to visualize gene expression in mouse brains.[20] This technique was developed concurrently to ours 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.

3.3 Parallel Coordinates and Scatter Plots

PC visualizations are most commonly used to search for patterns in non-spatial data. In its basic form, PC displays one 2D axis for each variable and represents samples as lines passing through those axes. To visualize uncertainty in this data, the most common technique is to use color.[21] However, this interferes with color's other common use in PC: line classification. Instead of color, Fua et al. draw lines with linear opacity falloff to represent variance in computed clusters.[22] We use a similar visualization of per-variable voxel variance, however we directly represent the normal distribution with an appropriately scaled Gaussian falloff that we interpolate between axes.

We use curved lines rather than straight lines to represent voxels. Curves are generally computed by enforcing either smooth derivative constraints at variable axes[23] or zero-derivative constraints at variable axes.[24] While the former lets viewers more easily distinguish lines incident on nearby points on an axis, it reduces cluster visibility and distorts line intersections between axes. Such intersections indicate linear relationships between variables,[25] an important feature we use in our visualization system. We therefore opt for zero-derivative constraints, which preserve the existence of intersections.

Scatter plots are a commonly used bivariate data visualization tool that display data samples as glyphs drawn on a Cartesian grid. Values of more than two variables are commonly conveyed by creating a matrix of scatter plots.[26,27] While such plots can display many of the same data relationships among pairs of variables, the parallel coordinates plot more directly shows relationships and population clusters among all variables by connecting the values on each variable axis.

3.4 Linked Visualization Systems

Many systems coordinate combinations of both abstract and spatial visualization techniques. XmdvTool[28] and GGobi[29] combine several abstract visualizations (PC, scatterplots, hierarchical views, etc.) of multivariate data sets. Akiba and Ma link interactions between PC, time histograms, and DVR.[30] Users can brush over interesting clusters of lines in the PC view or features in the histogram view and see the DVR update interactively. Systems like SimVis[31] and WEAVE[32] combine information visualization techniques (statistical representations, feature analysis, ND projections) with scientific visualizations. Our work uses a similar linked visualization approach to help radiologists understand relationships and values in MRS data.

4. Implemented Visualization System

The system we present to the radiologists has three components: a spatial SDDS visualization, a PC visualization that incorporates data variance, and a value extraction interface that uses a color-mapped anatomical slice plane.

4.1 Scaled Data-Driven Spheres

The SDDS technique distributes a separate set of spherical glyphs into 3D space for each metabolite. Bokinsky's 2D Data-Driven Spots map scalar value to glyph opacity.[19] As shown by Interrante et al., glyphs with varying opacities are extremely difficult to understand in depth.[15] In order to extend this technique into 3D, we therefore map data value to glyph radius instead, which has the added benefit that small glyphs in low-valued regions do not occlude high-valued regions. We use sphere glyphs because they are easy to interpret at any scale, even when small. Spheres placed along a regular grid exhibit strong aliasing and stacking effects, so we resample the scalar field on a randomly jittered version of the original sample grid, as shown in Figure 2. The glyphs for each scalar volume use separate jittered grid positions. We compute sphere radius as follows:

$$r_{max}=k\frac{s}{2n} \tag{1}$$

$$r=r_{max}\frac{v-v_{min}}{v_{max}-v_{min}} \tag{2}$$

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 overlapping. When sphere glyphs get too large, the viewer tends to lose a sense of the continuity in the data signal between voxels. For low resolution data sets such as MRS, Interpolating between data values with smaller glyphs solves the problem.

When displaying multiple scalar fields, we color each field of glyphs uniquely. Humans can quickly name and differentiate approximately 12 color values,[33,34] so this number serves as the theoretical upper limit of the number of simultaneously displayable scalar fields. Bokinsky found that viewers had no trouble visually separating nine fields of scaled and colored spots in a 2D Data-Driven Spots visualization. The inevitable increase in visual occlusion resulting from adding a new field limits the number of displayable scalar fields. Depending on the desired glyph size, sample density, and data sparsity, we have observed that images with more than five to six scalar fields become over-occluded and hard to interpret.

We combine SDDS with a 2D anatomical slice plane so that radiologists can correlate anatomical structure to metabolite values and relationships. Users have standard control over window and level values. We did not include a 3D visualization of the anatomy because SDDS produces a visually dense image; an additional 3D anatomy visualization would severely obstruct the spheres. An added benefit of the 2D slice plane is that it can be used as a partial occluder; when viewers are only interested in values near a particular z-slice of anatomy, the slice plane can be moved to occlude the color spheres beneath that slice.

Both stereo display and an interactive camera help viewers disambiguate depth in the renderings. As a result, they have a fuller understanding of the shape of the different scalar fields. Figure 2 shows a cross-eyed stereo image pair of SDDS to show the strength of the stereo effect. Note that even monovariate volume visualization of unfamiliar asymmetric structures is a difficult problem that benefits from stereo 3D. Stereo becomes even more important when viewing multivariate volume visualizations, regardless of the chosen technique.

4.2 Parallel Coordinates

SDDS helps users identify spatial relationships among metabolites and anatomy. We link the SDDS visualization to a PC visualization to help users more accurately quantify those relationships and select voxels that match. The PC visualization is modified from traditional PC in four ways. First, we use S-shaped cubic splines instead of straight lines. Second, we represent value uncertainty as Gaussian falloff with interpolated standard deviation. Finally, we incorporate two novel interaction techniques:

- **Linear Function Brushing**: Users can select curve patterns commonly visible in PC.
- **New Axis Construction**: Users can create new axes based on discovered relationships.

Perception research has shown that viewers can more easily follow curved lines than connected line segments.[33] This phenomenon is often explained via the Gestalt principle of "common fate." We therefore draw curved lines rather than straight segments as is traditionally done in PC. Cubic splines with smooth first derivatives are easy to follow and help viewers distinguish lines incident on the same position on an axis.[23] However, line clusters (both in slope and axis intersection) are easier to see when curves arrive perpendicular to the axes.[24] Such curves are S-shaped between axes. We opt for S-shaped curves because clusters are more important than individual curve behavior in our application.

4.2.1 Uncertainty—The MRS data for every metabolite in a voxel is represented by a scalar concentration and a standard deviation. Techniques we attempted for augmenting SDDS to incorporate uncertainty (such as opacity, with or without surface texture) result in confusing and cluttered visualizations. PC gives us the opportunity to represent this uncertainty. We extend the linear falloff technique used by Fua et al.[22] by using a Gaussian distribution that matches the voxel standard deviation. Because the standard deviation changes at each axis, we use the same zero-derivative spline computation used for the curves to interpolate standard deviation between axes. When the lines become extremely blurry, the original data values are completely lost. Although this accurately portrays the uncertainty in the data and is useful for showing general trends, it can interfere with the display of relationships. We therefore emphasize the center of each blurred curve (the distribution mean) by making it slightly brighter than the rest of the curve. Figure 4 shows the difference between blurred curves with and without mean emphasis. This brightness is a user-controlled parameter.

4.2.2 Linear Function Brushing—The patterns apparent with PC most often fall into two categories: clusters of lines passing through a small range of values or clusters of lines representing a linear relationship between two variables. For example, strong negative correlation results in a characteristic hourglass shape in the PC visualization. Figure 5a depicts some other patterns resulting from linear relationships. Our radiologist colleagues wish to use MRS to distinguish benign tissue from malignant tissue. Ideally, multiple sets of relationships would be visible in the PC plot. We therefore need interactions for selecting

curves satisfying a particular relationship, such as the linear relationships described above. We provide three mechanisms for selecting curves in the PC visualization: clustered curve selection, slope-based selection, and linear function-based selection.

To select a clustered set of curves, the user brushes over a region. A curve is selected if that region spans the majority of the integrated Gaussian. Once a curve has been selected, it is drawn in a different color and a corresponding wireframe cube appears around the voxel in the SDDS visualization. This lets users immediately understand the spatial context of their selection. However, selecting a set of curves can be difficult when multiple noisy relationships are present due to curve overlap. Angular brushing, described by Hauser et al., resolves this problem by selecting curves that have a range of slopes defined by the user.[35] Lines with similar slope signify a particular linear relationship; linear function brushing extends this to allow users to select the characteristic shapes of all linear relationships.

Users highlight more complex linear relationships by drawing two representative lines that describe a perceived visual pattern. For example, a user can outline both sides of the hourglass shape indicative of a negative correlation. We call this linear function brushing. We evaluate how well one curve matches the function represented by these curves as follows:

$$\left|\frac{a-a_0}{a_1-a_0}-\frac{b-b_0}{b_1-b_0}\right|<K \tag{3}$$

As shown in Figure 5a, a and b are the values of the curve being evaluated at two axes (A and B). a_0 and b_0 are the values of the first representative lines. a_1 and b_1 are the values of the second representative line. K is a threshold set by the user. The linear function described by these lines follows:

$$b=a*scale+shift \tag{4}$$

$$scale=\frac{b_1-b_0}{a_1-a_0} \tag{5}$$

$$shift=-\frac{a_0(b_1-b_0)}{a_1-a_0}+b_0 \tag{6}$$

We display the function described by the two representative lines for the user's reference. Note that when the axes are rescaled to individual variable ranges, angular brushing selects lines that indicate both a shift and a scale, similar to linear function brushing. While some parallel coordinates research has advocated forcing all axes to the same data range to preserve relationships, this can result in unusable columns for variables with very different ranges. Linear function brushing enables users to use visual pattern matching to discover relationships between individually scaled axes and informs them of the specific relationships those patterns represent.

4.2.3 New Axis Construction—Our radiologists find pair-wise relationships useful, however more complex relationships between variables also exist. Therefore, once a user has identified a meaningful combination of two variables, we let them create a new axis that represents that relationship. The new axis appears as a new set of spheres in the SDDS visualization. If the user finds this relationships using the brushing techniques described above, our system presents to them that relationship and allows them to create a new axis in one of two manners. The linear combination can be represented as either a difference ($y - mx - b$) or a ratio ($\frac{y-b}{mx}$). In the former case, curves that match the difference will be near zero, as shown in Figure 5b; in the latter, matching curves will be near one. The ratio form exacerbates differences for small values of x, so we additionally give users manual control over axis data range.

4.2.4 Other Features—Axis ordering is a common problem in PC visualizations. If a viewer sees two variables that may have an interesting relationship in SDDS, we let users simply drag and drop axes to swap the positions of two variables. Also, if the user is interested in only quantifying a linear relationship rather than selecting new voxels, we have implemented a query mode that provides this functionality. There are four selection operators that enable users to add to, subtract from, replace, and query the current selection (without selecting curves).

5. Use Case: Tumor Segmentation

We now show a use case of our visualization system applied to a data set with a tumor visible in the anatomical T1 slice plane. The tumor is visible in the slice plane because the patient has taken an intravenous contrast agent. In this data set there are four metabolites: choline (Cho), creatine (Cr), glutamine (Gln), and N-acetylaspartate (NAA). First, the user views the SDDS visualization to get a global sense of the spatial relationships between variables, as shown in Figures 6a and 2. In this image, several relationships are apparent. All of the metabolites appear depressed inside of the tumor, except for a small voxel in orange (Cr). The green (NAA) and orange (Cr) metabolites appear to be positively correlated outside of the tumor. Also, green (NAA) and orange (Cr) appear to be negatively correlated with yellow (Gln) outside of the tumor.

The user now seeks to confirm and quantify these hypotheses by looking at the PC visualization, shown in Figure 6d. First, the user orders the axes such that some of their hypotheses can be tested (Cho, Cr, Gln, NAA). This view shows a cluster of lines traveling from Cho to Cr in a linear function pattern. Using the linear function brush, the user draws two representative lines and selects voxels that match that pattern: $Cho \approx .9\ Cr - .05$, which is a strong positive correlation as shown in Figure 6b. Following the lines to the other two variables, it becomes clear that similar linear relationships exist in the other variables, and strong negative correlation appears to exist for the selected voxels between Gln and NAA. The user looks at the selected voxels in the 3D visualization and sees that this selection of voxels seems to match tissue outside of the tumor, shown in Figure 6e.

Returning to the PC visualization, the user seeks to confirm the positive correlation hypothesis between NAA and Cho and swaps NAA with Gln. Using query mode, the user draws the function that matches the previously selected voxels between the Cr and NAA columns and sees that they are indeed strongly correlated: $Cr \approx .4\ NAA$. The user now confirms the negative correlation hypothesis between NAA and Gln outside the tumor by applying a query brush a second time: $NAA \approx -.6\ Gln + .2$.

The user sees another cluster of curves with similar negative slopes between Cho and Cr. Choosing a second selection color, the user selects those curves with the angular brush, shown in Figure 6c. Note that the axes have been individually rescaled such that Cr has twice the range of Cho. Because of this, the negative slope selected actually represents a positive scale in data space: $Cho \approx .4\ Cr$. Returning to the 3D visualization shown in Figure 6f the user can see that the newly selected voxels are inside of the tumor. The selection of these two class of lines has indicated that the *Cho/Cr* ratio is smaller inside of the tumor than outside.

Four major insights have come from use of the visualization system:

- Outside the tumor, Cho and Cr are positively correlated: $Cho \approx .9\ Cr - .05$.
- Outside the tumor, Cr and NAA are positively correlated: $Cr \approx .4\ NAA$.
- Outside the tumor, Gln and NAA are negatively correlated: $NAA \approx -.6\ Gln + .2$.
- Inside the tumor, Cho/Cr is significantly smaller than outside: $Cho \approx .4\ Cr$.

The user can now compute a new axis to produce a tumor segmenter. Noticing that Cho, Cr, and NAA are all positively correlated to each other and negatively correlated with Gln outside of the tumor, the user can combine the above observations into a single function that will identify all such voxels:

$$Cho + .1\ Cr + .6\ Cr - .6\ Gln - .25 \approx 0 \tag{7}$$

The closer to zero the value of this function for a voxel, the more likely the voxel is to be outside of the tumor. This new axis appears as a new set of spheres in the SDDS visualization. The user can now use the 2D colored slice plane on the computed axis to extract specific values.

6. Conclusions and Future Work

The SDDS visualization successfully enables our radiologist colleagues to understand relationships between metabolites and anatomical features. While our visualization design was guided solely by radiologists, SDDS has a wider range of applicability than MRS. The SDDS visualization enables viewers to explore and analyze the relationships between the variables in a multivariate volume scalar field. We have also applied these techniques to multivariate chemical concentrations from the Virtual Cell project (http://www.nrcam.uchc.edu/) and multi-fluorophore confocal optical microscopy data sets. To our knowledge, SDDS is the only multivariate scalar volume visualization technique that has the potential to scale to 11 simultaneous channels of display. Currently, sphere scale is under user control, however future work on SDDS may include discovering what sphere scale leads to the greatest understanding.

In addition to SDDS, we investigated several other multivariate scalar volume visualization techniques. We presented isosurfaces, glyph-textured isosurfaces, and superquadric glyphs to our clients in addition to SDDS. The radiologists decided that SDDS most accurately enabled them to understand relationships among metabolite concentrations and anatomical structure. Opaque isosurfaces occluded each other, but adding transparency or texture produced images that were either too confusing or overly occluded. The surface-based techniques were able to identify specific structures, but relationships and values through the data volume were more difficult. Superquadric glyphs, a shape-based alternative glyph technique, produced shape representations for which the range of distinguishable values was

significantly different for different shape channels. Additionally, mapping semantically related variables (different metabolites) to semantically different channels (shape, color, scale, etc.) was confusing for our colleagues.

The PC visualization complements the SDDS visualization by discarding spatial relationships and focusing instead on clusters of voxels with similar behavior. We show an intuitive method for incorporating normally distributed uncertainty into PC visualizations. We will continue to investigate how to incorporate uncertainty into other information visualization techniques, such as scatter plots.

The brushing techniques we describe help users isolate specific voxel relationships via visual pattern matching without requiring the viewer to compute the relationships on their own. It is particularly useful for helping viewers identify patterns between axes that have individually rescaled data ranges. Rescaling all axes to the same range comes at the cost of sample resolution, and may not be possible for multi-modal data sets in which variables can have dramatically different ranges, and even units. This will be important for comparing MRS data to other imaging modalities. We will continue to investigate user-guided relationship discovery techniques. For example, user-identified relationships may be useful for seeding automatic clustering techniques that compute multivariate relationships for segmentation.

Our novel brushing techniques applied to PC enabled radiologists to identify relationships in MRS data. Gaussian curve representation provides a more accurate portrayal of trends and patterns in the presence of uncertainty. By linking PC with SDDS, radiologists can confirm hypotheses about relationships among metabolites and anatomy. As we improve both visualization and interaction with our system, we hope to continue to help radiologists diagnose patients with more confidence and accuracy.

Acknowledgments

Thanks go to our colleagues in the UNC Chapel Hill computer science department as well as the anonymous reviewers for their helpful comments and suggestions. This work was developed by the Center for Computer Integrated Systems for Microscopy and Manipulation (CISMM), an interdisciplinary science research group at UNC. CISMM is supported by the NIH National Institute of Biomedical Imaging and Bioengineering (NIH P41-EB002025).

References

1. Castillo, M. Neuroradiology. Lippincott Williams & Jenkins; 2002.
2. Soares DP, Law M. Magnetic resonance spectroscopy of the brain: review of metabolites and clinical applications. Clin Radiol Jan;2009 64:12–21. [PubMed: 19070693]
3. Provencher S. Estimation of metabolite concentrations from localized in vivo proton NMR spectra. Magn Reson Med Dec;1993 30:672–679. [PubMed: 8139448]
4. Maudsley AA, Darkazanli A, Alger JR, Hall LO, Schuff N, Studholme C, Yu Y, Ebel A, Frew A, Goldgof D, Gu Y, Pagare R, Rousseau F, Sivasankaran K, Soher BJ, Weber P, Young K, Zhu X. Comprehensive processing, display and analysis for in vivo MR spectroscopic imaging. NMR in Biomedicine 2006;19:492–503. [PubMed: 16763967]
5. Chang J, Thakur S, Perera G, Kowalski A, Huang W, Karimi S, Hunt M, Koutcher J, Fuks Z, Amols H, Narayana A. Image-fusion of MR spectroscopic images for treatment planning of gliomas. Medical Physics January;2006 33:32–40. [PubMed: 16485406]
6. Inselberg A. The plane with parallel coordinates. The Visual Computer 1985;1(4):69–91.
7. 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 2006:305–318.

8. Rheingans, P. Proceedings of IEEE Visualization '92. IEEE Computer Society; Los Alamitos, CA, USA: 1992. Color, change, and control for quantitative data display; p. 252-259.
9. Kniss, J.; Kindlmann, G.; Hansen, C. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. VIS '01: Proceedings of the conference on Visualization '01; Washington, DC, USA: IEEE Computer Society; 2001. p. 255-262.
10. Nattkemper TW. Multivariate image analysis in biomedicine. J of Biomedical Informatics 2004;37(5):380–391.
11. Broersen, A.; van Liere, R. EuroVis. Eurographics Association; 2005. Transfer functions for imaging spectroscopy data using principal component analysis; p. 117-123.
12. 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.
13. Sauber N, Theisel H, Seidel HP. 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]
14. Woodring J, Shen HW. Multi-variate, time varying, and comparative visualization with contextual cues. IEEE Transactions on Visualization and Computer Graphics 2006;12(5):909–916. [PubMed: 17080816]
15. Interrante V, Fuchs H, Pizer S. Conveying the 3d shape of smoothly curving transparent surfaces via texture. IEEE Transactions on Visualization and Computer Graphics April-June;1997 3:98–117.
16. Weigle C, Taylor RM. Visualizing intersecting surfaces with nested-surface techniques. IEEE Visualization 2005;64
17. Kindlmann G, Westin CF. Diffusion tensor visualization with glyph packing. IEEE Transactions on Visualization and Computer Graphics 2006;12(5):1329–1335. [PubMed: 17080869]
18. 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.
19. Bokinsky, A. PhD thesis. UNC; Chapel Hill: 2003. Multivariate Data Visualization with Data-Driven Spots.
20. 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]
21. Haroz, S.; Ma, KL.; Heitmann, K. Multiple uncertainties in time-variant cosmological particle data. Proceedings of IEEE Pacific Visualization Symposium; IEEE VGTC; Mar. 2008 p. 207-214.
22. Fua, YH.; Ward, MO.; Rundensteiner, EA. Hierarchical parallel coordinates for exploration of large datasets. VIS '99: Proceedings of the conference on Visualization '99; Los Alamitos, CA, USA: IEEE Computer Society Press; 1999. p. 43-50.
23. Graham, M.; Kennedy, J. Using curves to enhance parallel coordinate visualisations. Information Visualization, 2003 IV 2003 Proceedings Seventh International Conference on; Jul. 2003 p. 10-16.
24. Moustafa, R.; Wegman, E. Multivariate Continuous Data - Parallel Coordinates. Statistics and Computing, Springer; New York: 2006. p. 143-155.
25. Inselberg, A.; Dimsdale, B. Parallel coordinates: a tool for visualizing multi-dimensional geometry. VIS '90: Proceedings of the 1st conference on Visualization '90; Los Alamitos, CA, USA: IEEE Computer Society Press; 1990. p. 361-378.
26. Cleveland, WC.; McGill, ME. Dynamic Graphics for Statistics. CRC Press, Inc.; Boca Raton, FL, USA: 1988.
27. Elmqvist N, Dragicevic P, Fekete JD. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. IEEE Transactions on Visualization and Computer Graphics 2008;14(6):1141–1148. [PubMed: 18989008]
28. Ward, M. Xmdvtool: integrating multiple methods for visualizing multivariate data. Visualization, 1994., Visualization '94, Proceedings., IEEE Conference on; Oct. 1994 p. 326-333.
29. Swayne DF, Temple Lang D, Buja A, Cook D. GGobi: evolving from XGobi into an extensible framework for interactive data visualization. Computational Statistics & Data Analysis 2003;43:423–444.

30. Akiba, H.; Ma, KL. A tri-space visualization interface for analyzing time-varying multivariate volume data. Proceedings of Eurographics/IEEE VGTC Symposium on Visualization; May. 2007 p. 115-122.

31. Doleisch, H.; Gasser, M.; Hauser, H. Interactive feature specification for focus+context visualization of complex simulation data. VISSYM '03: Proceedings of the Symposium on Data Visualisation 2003; Aire-la-Ville, Switzerland, Switzerland: Eurographics Association; 2003. p. 239-248.

32. Gresh, DL.; Rogowitz, BE.; Winslow, RL.; Scollan, DF.; Yung, CK. Weave: a system for visually linking 3-d and statistical visualizations, applied to cardiac simulation and measurement data. VIS '00: Proceedings of the conference on Visualization '00; Los Alamitos, CA, USA: IEEE Computer Society Press; 2000. p. 489-492.

33. Ware, C. Information visualization: perception for design. Morgan Kaufmann Publishers Inc.; San Francisco, CA, USA: 2000.

34. Healey CH. Choosing effective colours for data visualization. Proceedings of IEEE Visualization '96 1996:263–270.

35. Hauser, H.; Ledermann, F.; Doleisch, H. Angular brushing of extended parallel coordinates. Proceedings of IEEE Symposium on Information Visualization; IEEE Computer Society Press; 2002. p. 127-130.

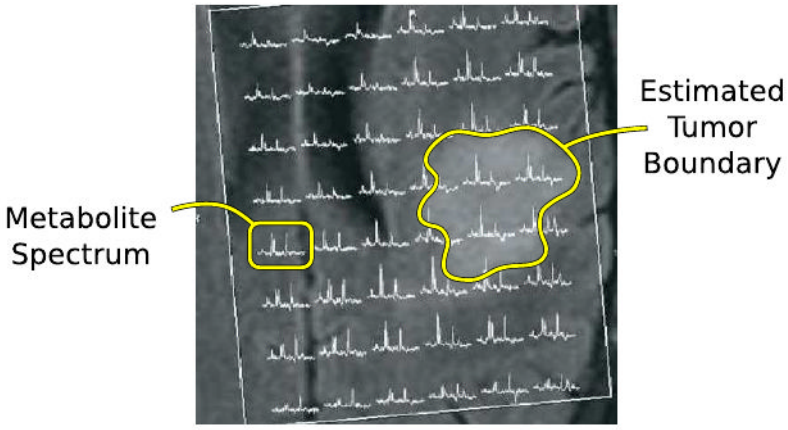

Figure 1.
Our radiologist colleagues currently visualize MRS data sets by overlaying metabolite spectra on anatomical MRI slices. This visualization is difficult to understand, even with special training.

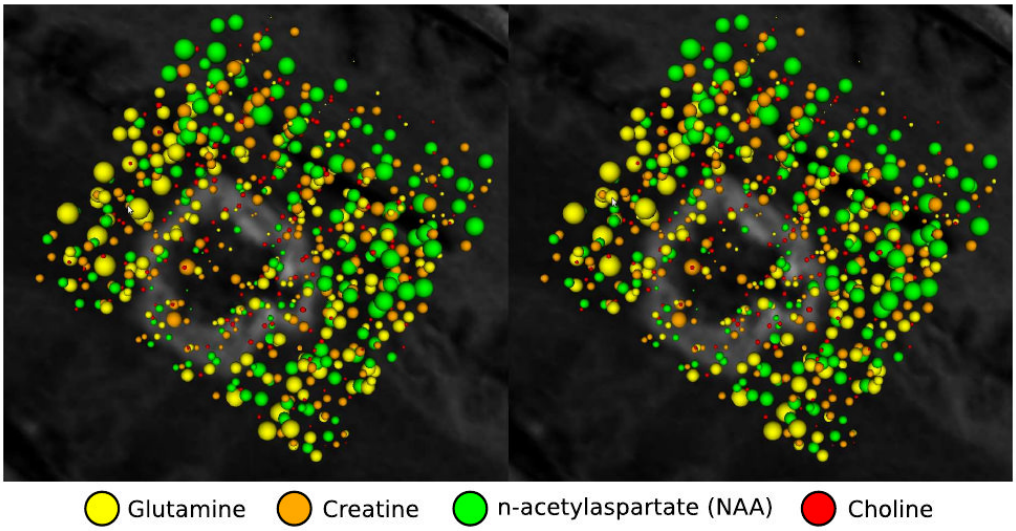

Figure 2.
A cross-eyed stereo pair of SDDS with four variables above a grayscale anatomical slice plane. The outline in the anatomical image is a constrast stain highlighting a potential tumor. Notice how all of the variables are low-valued within the tumor.

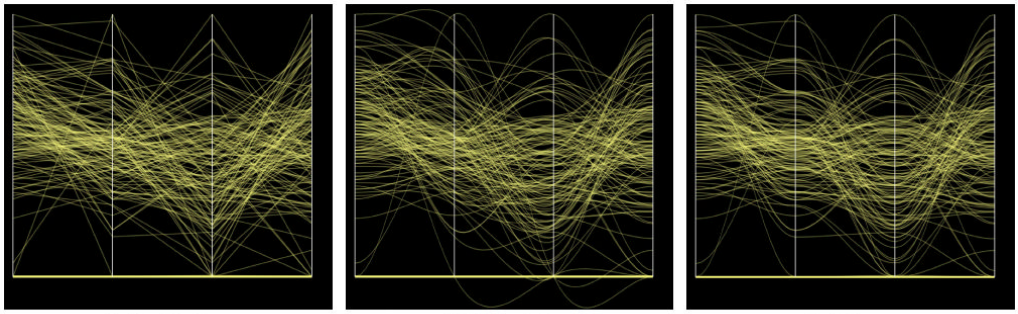

Figure 3.
Examples of three curve representations in PC visualizations. Straight lines are the traditional representation, but curves are easier to follow. Cubic splines help viewers track curves entering a small range of values. S-Curves emphasize clusters at the expense of tracking.

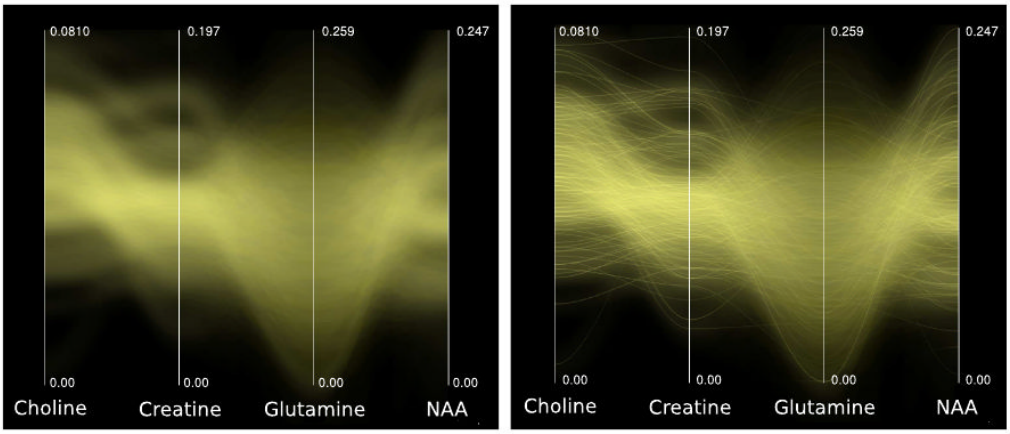

(a) Without Line Emphasis (b) With Line Emphasis

Figure 4.
Two ways of representing uncertainty, with and without emphasis of the original data point value. Center line emphasis (4b) highlights mean metabolite concentrations at the cost of implying a more certain value. The degree of emphasis is a user-controlled parameter.

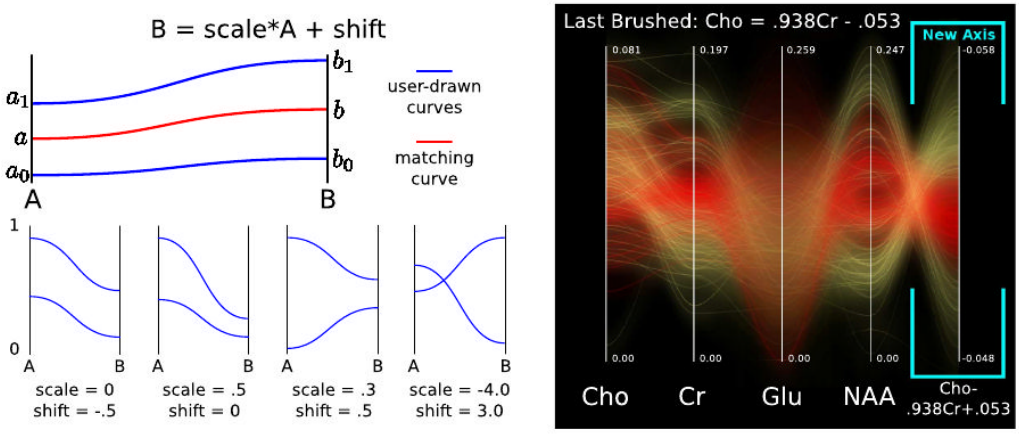

Figure 5.

Left: diagram of variables used in the equations in Section 4.2.2 and example visual patterns that emerge when linear relationships among variables. Using linear function brushing, the user draws two representative curves to select all curves that match a particular linear function. Right: once selected, a new column representing the selected function can be constructed to compare that relationship to other variables.

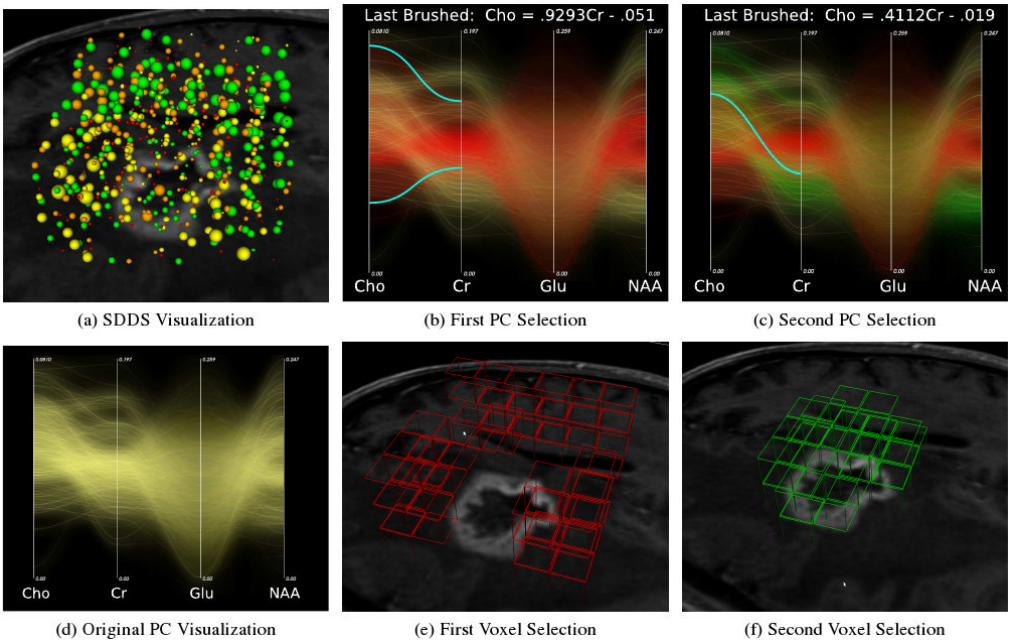

Figure 6.
Use case images. Left: The user confirms hypotheses gathered from the SDDS visualization by selecting visual curve patterns in the PC visualization. Center: The first pattern selects voxels predominantly outside the tumor. Right: The second pattern selects voxels predominantly inside the tumor.