SpectraMosaic: An Exploratory Tool for the Interactive Visual Analysis of Magnetic Resonance Spectroscopy Data

L. Garrison[1,3], J. Vašíček[2], R. Grüner[3,4†], N.N. Smit[1,3†] and S. Bruckner[1,3†]

[1]Dept. of Informatics, Univ. of Bergen, Norway [2]Dept. of Informatics, Masaryk University, Brno, Czech Republic [3]Mohn Medical Imaging and Visualization Centre, Dept. of Radiology, Haukeland Univ. Hospital, Bergen, Norway [4]Dept. of Physics and Technology, Univ. of Bergen, Norway
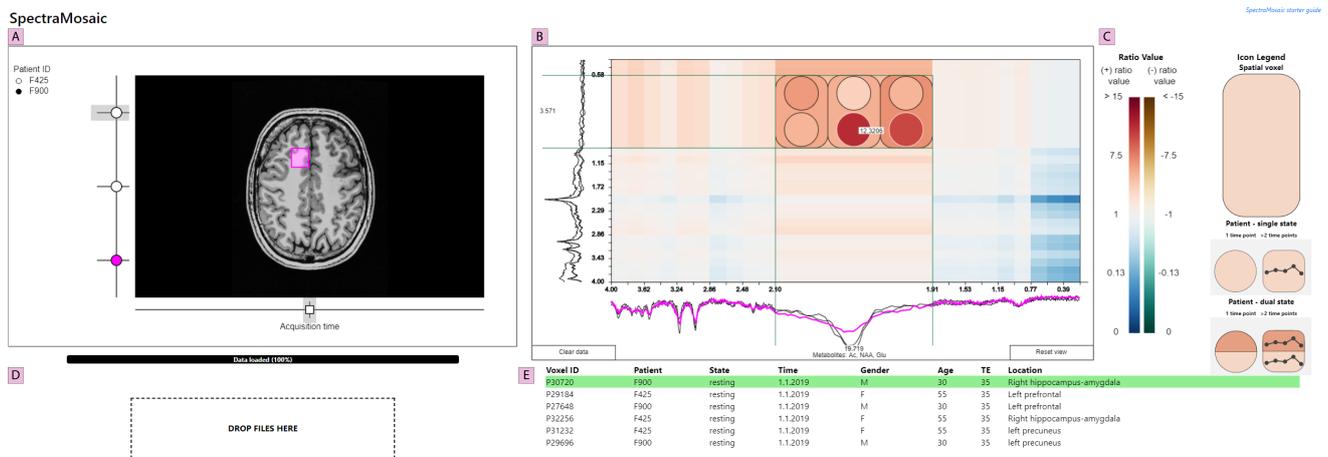
Figure 1: *SpectraMosaic supports exploration and comparison of magnetic resonance spectroscopy (MRS) metabolite ratios with an anatomical view (A), a spectral heatmap panel (B), a legend view (C), data loader (D) and metadata table (E).*

Abstract
Magnetic resonance spectroscopy (MRS) allows for assessment of tissue metabolite characteristics used often for early detection and treatment evaluation of brain-related pathologies. However, a steep learning curve for metabolite interpretation, paired with limited visualization tools, have constrained the more widespread adoption of MRS in clinical practice. In this design study, we collaborated with domain experts to design a novel visualization tool for the exploration of tissue metabolite concentration ratios in MRS clinical and research studies. We present a data and task analysis for this domain, with categorization of MRS data attributes into tiers of visual priority. We furthermore introduce a novel set of visual encodings for these attributes. Our result is SpectraMosaic (Figure 1), an interactive insight-generation tool for rapid exploration and comparison of metabolite ratios. We validate our approach with two case studies from MR spectroscopy experts, providing early qualitative evidence of the efficacy of the system and affording deeper insights into these complex data.

CCS Concepts
• Human-centered computing → Scientific visualization; Information visualization; User centered design;

1. Introduction

Magnetic resonance spectroscopy (MRS) is an *in vivo* non-invasive biochemical imaging technique utilized for tissue metabolite characterization. Metabolites are chemical compounds that are the end product of metabolism, a process necessary to maintain life. In medical research, MRS has been used in psychiatric and neuro-

logical studies, including those for tumor tissue differentiation, Parkinson's disease, stroke causality, neonate oxidation status, and schizophrenia. As it is extremely sensitive to subtle local tissue composition changes, clinical researchers have begun to explore MRS as a tool for tracking metabolite variation over time, space, and individuals [VDG10]. Although recent technology improve-
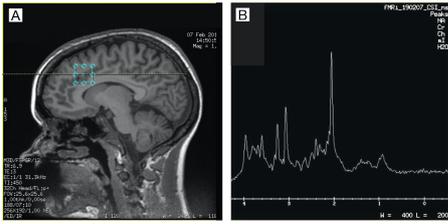
Figure 2: *Typical visual output of MRS from a clinical imaging workstation with spatial context for the spectral voxel in a T1-weighted high resolution MRI slice view (A) and spectral peaks quantified within that voxel (B).*

ments in MRS acquisition have produced higher data resolution and improved signal-to-noise ratio (SNR) [VKV*16], visualization tools for interpretation and comparison are still limited. The current visualization standard in most medical imaging application systems is placement of a voxel cube with a spectral graph juxtaposed with or superimposed onto a high resolution anatomical image (Figure 2), occasionally with a single metabolite heatmap overlay [NLK*14]. A spectral graph is visually similar to a function graph (Figure 2B), where peaks in the spectrum represent metabolites. The x-axis of the graph indicates chemical shift values used to describe metabolite characteristics while the y-axis indicates signal intensity. The abstract nature of this visualization method creates a steep learning curve to spectral interpretation, often requiring rote memorization to derive insights.

MRS data is relatively unexplored in visualization research. Approaches have focused on different specific tasks in MRS multi-voxel spectroscopy data analysis than those explored in our application [FKLT10a,FKLT10b,FLKT09,NRS*14]. We have designed SpectraMosaic to provide an simple entry point to spectral interpretation while offering the flexibility to explore and form new insights through multiple facets of the data. The goal of this work is to show the range of different metabolite concentration ratios at different dimensional tiers, represented through basic visual encodings, to aid in MRS data interpretation and analysis. Our main contributions of this design study include:

- MRS data and task analysis with subsequent task abstraction, guided by domain expert needs
- Proposed system of visual encodings for MRS data
- Iterative development of SpectraMosaic, a novel MRS visualization tool, in close coordination with target end users
- Validation of our design approach with two case studies illustrating the utility and value of SpectraMosaic in addressing user needs.

Flexible and simple to use, SpectraMosaic is the first MRS application to allow interpretation of processed metabolic spectra over four dimensions: space, time, state, and cohort size.

2. Related Work

Table Visualization First introduced by Bertin [BBW83], numerous solutions leveraging small related graphics series have

been developed as a method for visualizing heterogeneous, multivariate, multidimensional data. SpectraMosaic draws inspiration from this concept but layers visualizations with dimensional priority, with secondary dimensions presented inside the first. PivotTable, subsequently trademarked by Microsoft and extended by Polaris [STH02], allowed exploration and analysis of multidimensional data with the flexibility to modify visual encodings, graphics, and table configuration for visualization. Our approach is related in that we allow on-the-fly reconfiguration of our table matrix inputs, but we extend this concept further by including a second layer of nested visual encodings as inspired by Atom [PDFE17], a grammar for unit visualizations where individual data items are represented by unique visual marks (units) in a visual encoding system.

Complex Heterogeneous Data Visualization InSpectr [AFK*14] utilizes multiple linked views from multimodal data sources (x-ray computed tomography and x-ray fluoroscopy) to provide insights into composition of a sample. Comparative visualization techniques as described by Gleicher et al. [GAW*11] are realized with magic lenses to show relative element distributions. SpectraMosaic similarly combines multimodal imaging techniques, but is applied to a medical domain and focuses on nesting spatial distribution data inside abstracted data. Isosurface similarity maps defined by Bruckner and Möller [BM10] were applied to spectra in Spectral Similarity Maps in an extension of the Inspectr framework, where correlations between spectra are shown as an intensity map. We adopt a similar concept in our tool, but rather than mapping energy correlation we instead map ratio strength between metabolites. Multeesum is a tool developed for visualization of gene expression over time within a spatial context and permits similarity measurements of gene expression and localization across species [MMDP10]. Of particular note is Multeesum's approach to presenting time-varying gene expression data in a small multiples view, which was first established in Pathline [MWS*10]. We are inspired by this handling and display of time series data in our visual encoding of time data.

MRS Visualization Research Prior visualization approaches for MRS data have been limited to the analysis and visualization of a subset of metabolites at a time. SDDS (scale driven data spheres) adopted by Feng et al. [FLKT09] provide a 3D representation of metabolites within a voxel. This application was later extended to include scatter- and parallel coordinate plots with limited metabolite types for comparison [FKLT10b]. SpectraMosaic remains in the abstract visualization space, but allows for comparison of all potential metabolite ratios within a given spectral set. Nunes et al. [NRS*14] presented a visual analysis framework combining ComVis [MFGH08] and MITK [WVW*05] where MRS data are plotted in linked histogram or scatterplot views with other molecular or anatomical data with different brushes to define biological target volumes. However, this work was developed specifically for radiotherapy treatment visualization, where retention of spectra was not the focus of the application. It also supported only simple comparisons of metabolite ratios. SpectraMosaic extends this flexibility of metabolite ratio calculations, and displays additional MRS data attributes (temporal, spatial, and individual) in a focus+context visual representation. Marino and Kaufman [MK11] implemented direct volume rendering (DVR) to represent male prostrate anatomy from MRI data combined with PET and MRS in

prostate tumor delineation. Their technique utilized color mapping and layering to represent different imaging modalities combined with a score qualification metric to validate the data. However, this application was focused on a single metabolite ratio, and could only present an individual in a single time slice.

SpectraMosaic uniquely leverages abstract visualization methods in a layered discovery approach to medical spectroscopy data. It is the first of its kind to input processed spectral curves directly into a visual analysis tool for deeper exploration of metabolite attributes and relationships.

3. Background

MRS works by detecting non-zero atomic nuclear spins, where energy is released as the nuclei return to their initial energy state after a radiofrequency excitation. This release of energy separates the MRS signal peaks in the spectral readout, because spectral peak position is determined by the degree of shielding provided by the electron cloud surrounding the atomic nucleus. This property is known as chemical shift, and is the basic premise enabling MRS [UBA16]. The area under each peak is proportional to the concentration of nuclei giving rise to this peak; a given metabolite may have single or multiple nucleic peaks that additively comprise its voxel sample concentration. Although MRS can be performed on any atomic nucleus with non-zero spin [VDG10], 1H (proton) MRS is used most often in clinical routine, capable of detecting metabolites in concentrations 10,000 times lower than that of fat or water as imaged in conventional magnetic resonance imaging (MRI). Additionally, proton MRS can be performed using a conventional MRI system, requiring no extra hardware or machinery. Signal acquisition techniques include single voxel spectroscopy (SVS) or multivoxel spectroscopy (CSI). CSI, essentially a slab of multiple smaller single voxels, covers a much larger area than SVS but the trade-off is that the signal-to-noise ratio (SNR) is often much lower than in SVS. Since SVS acquisition techniques afford more detailed spectra for analysis we focus the remainder of our work on this technique.

Following signal acquisition, a series of processing steps must be taken to quantify the MR spectra and produce the spectral output graph as seen in Figure 2. LCModel [Pro01] is a widely-used commercial tool used for processing MRS data, while open source tools such as jMRUI [SDCA*09], TARQUIN [WRK*11], SIVIC [CON13], OXSA [PCB*17], and Gannet [MMO*14] offer open-source solutions to spectral quantification with rudimentary visualization options. The typical output from these tools is a spectral plot of signal intensity against frequency – this can be representative of a single time slice as in a longitudinal study, or as a time-resolved functional study. Individual metabolite concentrations vary by their peak integral(s). Because acquisitions can vary considerably depending on the patient and sample location, results are often reported as ratios between metabolites, and for group analysis significance thresholding is applied. MRS data are typically visualized relative to a high resolution structural image to provide positional context [LSBP18]. These tools, as well as MRS standard clinical imaging workstation tools, display spectral curves placed adjacent to the voxel position image (see Figure 2), or with single-metabolite spectral color maps overlaid onto the complementary structural image. SpectraMosaic does not aim to compete

with these spectral processing tools, but rather provide a new way to visualize spectral data following the spectral processing step.

4. Data and Task Analysis

We worked iteratively with domain experts through a series of individual and group interviews to delineate and abstract this problem domain into a set of core user task and design requirements. Our collaborators, one of whom is a co-author of this paper, include 2 MD/PhDs, 3 PhD researchers, and one engineer, all specializing in MR spectroscopy research and working primarily with SVS data in an academic hospital setting.

Data Collection The majority of spectral acquisitions are static, meaning a single session captures a single moment in time of metabolite concentrations in a tissue region. Less common is a functional approach, which considers the variation of metabolite concentrations within an acquisition session period. In this approach the subject can also be asked to perform tasks, such as tapping fingers during the acquisition, and alternately resting – these time- and state-resolved acquisitions consider the same four attributes as in a static acquisition, but represent different scenarios and are important for the researcher to be aware of in their analysis. The signal is then output into a P-file, a file format proprietary to GE Healthcare MR systems. Subsequent processing and quantification of these raw P-files in TARQUIN or LCModel generate the following information for visual analysis: an averaged raw metabolite spectrum, a spectrum baseline, a model fit, and a list of the peak frequencies of all metabolites. Because MRS data is patient data, special care must be taken to anonymize these data for use in analysis applications. Following data collection researchers then perform a series of tasks that can be grouped into the following core tasks as defined by Brehmer & Munzner [BM13]: data discovery, production, search, and querying.

Tasks We use Brehmer & Munzner's multi-level typology of abstract visualization tasks [BM13] in our MR spectroscopy task analysis. We follow each of our four abstracted tasks with the concrete MRS domain tasks falling therein.

The first set of tasks relate to data consumption to **discover and verify key information about the source data (T1)**. MRS spectra, anatomical reference images, and associated anonymized patient data are first imported into an analysis tool. To analyze spectra and verify hypotheses it is critical to correlate spectral peaks with their corresponding metabolites. Although metabolites have constant chemical shift values, variations in acquisition parameter settings may change the shapes of, or hide, metabolite peaks [SR13]. A second core discovery task requires matching the voxel sample location to the spectral output as a means to contextualize and understand normal versus aberrant spectral outputs.

Secondly, **data are derived to produce new computed data elements from raw spectral input for analysis (T2)**. Because metabolic spectra vary considerably between and within individuals, peak integral ratios calculated from the raw spectrum are the gold standard for understanding variability. The use of ratios is a core critical task for any MRS application for two reasons, (1) as

a method to correct for homogeneity issues across samples and (2) as a basis for subsequent tissue composition analysis.

A third task set involves **selection and filtering of MRS input or derived data for exploration and comparative analysis (T3)**. Our collaborators often wish to select a subset of spectra for further analysis – they may wish to look only at the variation at a single time point in a longitudinal population study, or to examine gender variation within a study. Performing tasks of this type means the user must search MRS voxel and structural image data to lookup subsets of interest using selection, filtering, and aggregation tools. This brings us to a related search and query task: discovery of outlier metabolite concentration ratios in a given spectral output. An outlier is often of interest because it could represent an aberrant metabolite concentration that may be indicative of a pathology. Outliers are also useful in quality assurance; unexpected or high numbers of outliers may indicate a need to rerun the acquisition or cull a particular spectrum from study analysis [SR13].

A final core discovery task involves data comparison and summarization, where researchers **search metabolite ratios via browsing and exploration to identify, compare, and summarize spectral ratios across key dimensions of interest via selection, filtering, and navigation (T4)**. For example, comparison becomes particularly useful in brain lesion mapping as researchers try to understand how biochemical composition of the lesion compares to the healthy side of the brain. Researchers also wish to develop deeper insights into their data for how spectral metabolite ratios vary spatially, temporally, individually, and between active or resting brain states. Each of these attributes for comparison represent opportunities to visualize outliers, offering new windows for understanding sources of variation and patterns within a study. For example, in the area of psychiatric research it is helpful for our collaborators to see how key neurotransmitters such as glutathione, GABA, and glutamate vary across individuals, brain state, and in response to treatment. It can be useful to investigate the changing concentrations of these metabolites over variable time scales and over different locations in the brain to understand the complexity in the data beyond statistical mean measurements.

MRS Data Tiers We can break MRS attributes into three tiers, ordered by visualization priority: **Tier 1**: Original spectral graph data, **Tier 2**: Derived spectral data, and **Tier 3**: Spectral metadata. The first tier is of primary importance and comprises the original spectral graph that is critical for overview and quality assessment. Our second tier is comprised of the metabolite concentration ratios we calculate and then filter through the following dimensions: time, space, patient, and brain state. Time indicates either the number of separate spectral acquisitions performed on an individual over a study period, as in a longitudinal study, or can indicate recorded metabolite values within an acquisition session, as in a functional MRS study. Space indicates the voxel sample position within the brain. Patient refers to the individuals included in analysis. Finally, brain state indicates if the subject was in an "active" (task-explicit) state or "resting" (task-negative) state during signal acquisition. A functional MRS approach may record both states in a single session. Tier 3 comprises metadata important for context and filtering that is unnecessary to include as explicit encodings in the visualization. Gender, age, voxel location, and acquisition settings comprise

Case	patient	voxel	time pt	state	encoding	Case	patient	voxel	time pt	state	encoding
1	single	single	single	single		9	multiple	single	single	single	
2	single	single	single	dual		10	multiple	single	single	dual	
3	single	single	multiple	single		11	multiple	single	multiple	single	
4	single	single	multiple	dual		12	multiple	single	multiple	dual	
5	single	multiple	single	single		13	multiple	multiple	single	single	
6	single	multiple	single	dual		14	multiple	multiple	single	dual	
7	single	multiple	multiple	single		15	multiple	multiple	multiple	single	
8	single	multiple	multiple	dual		16	multiple	multiple	multiple	dual	

Figure 3: *Breakdown of all 16 possible permutations of MRS case study scenarios with sample encodings for each case. Key attributes fall under four main categories: patients (single individual, multiple individuals), voxels sampled (single voxel, multiple voxels), acquisition runs (single time point, multiple time points), and acquisition state (single state, or dual: active, resting).*

other important patient attributes to track because the shape of the spectrum can vary considerably with these factors [XV10, SR13].

Design Requirements Our colleagues often switch between hospital workstations while accessing sensitive patient data. A web application permits seamless workflow as researchers move between different stations and avoids downloading third party software to machines in a setting where security and privacy is paramount **(R1)**. Sensitive patient data must also be properly anonymized prior to loading into the visualization tool, achieved by culling identifying attributes from the source scan files to produce de-identified data files **(R2)**. To further support **T1, discover and verify key source data information**, we can map chemical shift numbers to peaks in satisfaction of T1 as a third design requirement **(R3)**, since chemical shift values are constant for all proton metabolite spectra. Lastly, we visually link our three data tiers in support of T1 to form our our fourth design requirement **(R4)**: maintenance of data linkage between voxel sample location information, spectral output, and patient-specific information.

Since user analysis tasks center around calculation and comparison of peak concentration ratios (T2-T4), this application must also support individual and aggregate concentration ratio calculations for any permutation of peaks in a spectral set **(R5)**. A layered design approach combining an aggregate ratio overview with nested individual ratio dimension information forms the basis of our final design requirement **(R6)**.

5. SpectraMosaic Visual Encodings

Because spatial resolution of metabolic spectra is low, we opted for an abstract heatmap matrix presentation of these data, as seen

in Figure 1B. Given that we wish to visualize metabolite concentrations as ratios, we place spectra (tier 1 MRS data) for comparison perpendicularly along an x- and y- axis to provide the inputs for a metabolite comparison matrix between the two axes. Since each spectrum typically has at most twenty interesting metabolite peaks [PODA13], we can discretize each continuous spectrum into 20 tile blocks – this allows us to produce a 20 by 20 cell table, where each cell corresponds to the ratio of averaged spectral tile integrals along the x versus y or y versus x axes (Figure 1B). We can then plot additional interesting information relevant to each spectral ratio (tier 2 data) inside each heat map tile. To provide the user with a range of hues for rapid visual inspection of areas of major difference we chose diverging colormap sets from ColorBrewer [HB03]. We discuss our color choices in more detail in Section 6.

Within each heatmap unit cell we define a nesting structure for spectral concentration ratios as inspired by dimensional stacking visualization techniques pioneered in Xmdvtool and N-land by Ward et al. [War94, WLT94]. We apply our visual system to each of the sixteen potential case scenarios in Figure 3. Each MRS attribute receives a consistent visual encoding in the form a simple glyph. We represent voxels as pill box glyphs, and in each unit cell evenly divide the space vertically. Patients are presented as filled disks when only shown in a single time acquisition (e.g., case 1), but expand into rounded squares when time series data are incorporated (e.g., case 3). This shape change is to permit a time spark line to move evenly across a space, rather than extend past or be cropped out by the rounded disk shape. Disks scale automatically to fill space optimally within their voxel frame to maximize pixel screen space. In instances where different brain states (active versus resting) are analyzed we break the patient disk/square horizontally in half. Finally, we encode time steps as points connected via a spark line, inspired by Meyer et al. in their work, Pathline [MWS*10]. This spark line is nested into the relevant glyph: if a multi time step series were captured in a study analyzing different brain states, the spark line is placed within each state half-moon "state" glyph (e.g., case 4). If analysis is only for a single state, the spark line nests inside the patient square glyph (e.g., cases 3, 11), or inside voxel glyphs if analysis is for a single patient (e.g., case 7). The remaining 16 cases comprise different permutations of these patient, voxel, state, and time arrangements.

6. SpectraMosaic

We now present SpectraMosaic, an exploratory tool for interactive visual analysis of complex heterogeneous MRS data. This design was guided by our data and task analysis outlined in Section 4 and by our visual encoding system outlined in Section 5.

SpectraMosaic Calculations At its core SpectraMosaic presents a visual comparison of single and aggregated metabolite ratios in a given MRS spectra collection based on data in the first and second tiers of our visual encoding system. The third tier of information is included in a header file to the spectral data and displayed as a table below the spectral heatmap. Although the general concept of peak integral ratio calculations are standard in MRS analysis, our layered approach described below is novel for MRS visualization.

We first normalize spectral input data along each axis to joint

maximum and minimum values. These values update with each modification to the spectral collection. Throughout these calculations we maintain connection between our header and data files so that semantic linkages are preserved (R4). Our image panel data are stored hierarchically, with relevant attributes such as TE setting, patient age and gender, as well as timepoint, state, and voxel ID. Once loaded into the spectral heatmap view (Figure 1, panel B) this hierarchy is flattened and the data are split into separate voxel arrays, one for each axis. For each MRS spectral voxel we store a label with including values for patient ID, voxel ID, state, time, scale in ppm. Each spectrum is divided into 1024 normalized samples – the vectors of these samples are also stored with each spectral voxel and used in integral calculations. These data values are summed into tile values, which are stored with the voxel and subsequently used to calculate spectral tile integral values then used for ratio calculations, approximated by the stored sum value within a tile region. We store 20 tile values per axis; these tiles can be resized by the user on-the-fly for customized analysis. Tile integral values are assigned either a positive or negative sign, as determined by the position of the corresponding metabolite peak(s) as either above or below the baseline. Peaks above the baseline, where the baseline is calculated as a pre-processing step in generating the spectral model fit, are assigned a positive sign in their integral calculations. Peaks falling below the baseline are correspondingly assigned a negative sign. These negative peaks are typically clinically significant as either pathological indicators or as indicators for changes in acquisition settings or techniques [SR13], and are therefore important to recognize in the visualization. Following tile integral calculations, all ratios are calculated as the average tile integral of x divided by the average tile integral of y.

We apply a diverging colormap to our normalized spectral table cell calculations, where cell color is determined by the ratio between the average tile integral on each axis. Our diverging color mapping function uses two maps; red-blue and gold-green. In our subsequent metabolite ratio calculations, red-blue is used if the calculated ratio is a positive value, while gold-green is used if the calculated ratio is a negative value. We selected these two maps from ColorBrewer [HB03]. Red-blue color maps are used frequently in scientific visualization and are a familiar sight to our collaborators, while green-gold shares an analogous color space and thus avoids generating too many disparate colors in the matrix space and overwhelming the viewer. In instances where our positive ratio calculations yield a value less than 1, we invert the ratio and switch the sign. Red indicates a higher ratio value while darker blue indicates a lower ratio value. Similarly, negative ratios using the green-gold map that are higher than -1 also are inverted with a sign switch to achieve the full range of divergent coloring. High ratios encoded as red or gold in hue mean that the metabolite(s) identified on the x-axis of the table matrix are present in greater concentration than the metabolite(s) on the y-axis. Lower ratios encoded in blue or green hues depending on the ratio sign indicate a greater concentration of the metabolite(s) initially identified on the y-axis of the matrix. Lastly, for the visualization we drop all values by 1 so the diagonal of the heatmap matrix is 0, rather than 1, to achieve a cleanly divergent color mapping structure in both maps.

On cell expansion we create a new data hierarchy that reflects the hierarchy of elements inside the cell: data from the spectral tile seg-
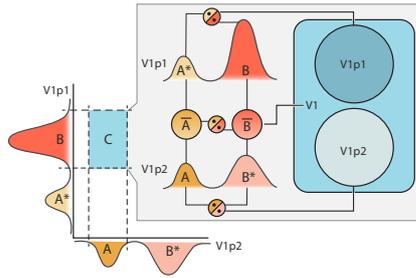
Figure 4: *Workflow image demonstrating the calculations necessary for a single-voxel two-patient case scenario (as in Figure 3, case 9). The overview spectral heatmap shows a single-voxel acquisition with the metabolite spectrum of patient 1 (V1p1) along the y-axis with the spectrum of patient 2 (V1p2) along the x-axis. The ratio of V1p2 peak A vs V1p1 peak B is shown in cell C, with dashed lines to indicate tile areas used in the ratio calculation for this cell. Subsequent expansion of cell C shows the average ratio in voxel 1 of peak A to B for both patients as a pill box, while the two nested disks represent peak A versus B for V1p1 and for V1p2, respectively. The darker blue coloring of V1p1 indicates a lower ratio value than V1p2.*

ment along both x- and y-axes are flattened into a single array and ordered first by voxel, then patient, followed by state, then lastly by time, for each present attribute. We then count the ratio for each voxel (average of all patients for this voxel), each patient (average of all states for given patient in a given voxel), each state (average of all timepoints for a given state of a patient from a voxel location) and each time point. Ratios are counted between tile regions of the same data for each visual tier. An example of our calculation workflow is demonstrated in Figure 4. In this scenario the user is interested in comparing the concentration ratio of peak A to peak B in a voxel sampled from the same brain region in two different patients. The spectral heatmap overview tile C shows the average metabolite concentration ratio of peak A (patient 1) to peak B (patient 2). This affords the researcher an opportunity to compare a metabolite peak from one patient against a different metabolite peak in a different patient, which is potentially useful in reference comparisons. Dashed lines leading from the x- and y-axes respectively indicate the spectral tile integrals used in the ratio calculation for cell C. Subsequent expansion of cell C shows a different set of calculations related to A versus B: the average ratio in voxel 1 of A to B for both patients as a rounded rectangle, while the two nested discs represent A versus B for patient 1 and for patient 2, respectively. The darker blue coloring for V1p1 shows that the ratio of A to B for patient 1 is much lower than patient 2. Researchers can compare the ratios of these unit cell visualizations against the cell ratio visualized in C as a further assessment of inter- and intra-patient metabolite variation.

Interaction and Workflow In designing SpectraMosaic we chose a multi-panel coordinated view (Figure 1); the panel A provides for selection of spectral voxels for analysis while panel B serves as the spectral exploration interface. This panel is a table matrix where spectra are placed along the x- and y-axes based on user selection

and allocation from the left panel. We use a pre-processing stack that anonymizes (R2) and extracts the raw data to a custom file format we use containing the relevant imaging and spectral data. Our reason for this is that, unlike other imaging modalities, DICOM formatting is not standardized for spectroscopy. Instead, medical imaging companies (e.g., GE, Siemens) each have different proprietary formats for their MR machinery. Study datasets are loaded into the application using a drag and drop window feature (Figure 1D)– these data remain linked semantically in the application to ensure continuity between voxel location, spectral output, and patient-specific information (R4). Spectral information is stored on the client while the anatomical image with voxel sample information is displayed in the voxel selection panel (Figure 1A). This panel consists of a set of images which can be navigated by patient along the vertical axis or by acquisition point (time) along the horizontal axis. In each anatomical image, a fuchsia square indicates the site in the brain where the spectral voxel was acquired. A vertical axis to the left of the anatomical image set shows small filled disks; each disk indicates the MRS acquisition type used for each image. Using the standard CPK coloring scheme for atomic elements seen in standard molecular visualization tools such as Rasmol [SMW95] and JMol [Her06], we represent a 1H collection with a white-filled disk. A light gray bar behind the disks show the image the user is actively viewing, while the disk becomes filled in fuchsia to indicate image linkage to a spectrum that is selected in the spectral heatmap panel (Figure 1B). Researchers may use this axis to traverse the sample stack to verify homogeneity of voxel position and voxel-image registration for quality assurance.

Data are loaded from the left voxel selection panel (Figure 1A) to the right spectral heatmap panel (Figure 1B) by clicking and dragging voxels from the left to right panel. Users can load as many spectra as they like, in any combination, onto each axis for more flexible analysis. The spectral map overview (Figure 1) dominates the right panel of the interface and is the main portal for metabolite comparison. Hovering with the mouse over any heat map unit cell shows a tooltip of the metabolites (R3) analyzed within that particular unit cell, as well as the tile region integral values for each axis and the mean metabolite concentration ratio for that cell (R5). Users may interactively expand the width and height of an interesting unit cell by clicking within the cell, as inspired by Bertifier [PDF14] for more detailed inspection (R5, R6). The expanded cell containing tier 2 attribute glyphs may be scaled to fill the entire space of the matrix in instances where larger cohorts are being analyzed, although MRS studies by our collaborators generally include only a handful of patients. Hovering over any of the unit cell elements correspondingly highlight the associated spectral graphs and image voxel slices in fuchsia (Figure 1A,B) as well as table rows in green below the heatmap visualization (Figure 1E) to correlate all visualization elements (R4). Figure 4 demonstrates unit cell expansion to compare detail metabolite ratios of individual elements (tier 2 data) contributing to the aggregate ratio color encoding. These unit elements scale automatically and uniformly for optimal space-filling in the cell. The background of the cell remains visible behind individual ratio elements for all expansions to preserve context of the aggregated value in navigation. A legend at the far right (Figure 1C) indicates the mapping between color and metabolite ratio,

as well as representing each glyph with its corresponding tier 2 MRS attribute.

7. Implementation

SpectraMosaic is a web-based application implemented with HTML, CSS, and p5.js for visualization of graphs and images in a Bootstrap framework (R1). Assets are stored on the client and fetched on-demand. Both left and right panels in the interface are drawn as p5 instances. We use the HTML Drag and Drop API to load data as a single root directory containing header and patient data; individual voxel data information are nested within each patient directory. Image data are read and converted by PNGReader to a p5 image displayed in the left panel. Data are shown in this panel based on patient, voxel, timepoint, and state selection. Once a voxel has been selected these data are stored as a global variable, which can then be accessed by the right spectral heatmap panel. Each global variable, once dragged to the x- or y- spectral heatmap axis, is then stored with that axis. When data have been added to both axes the heatmap is drawn as a 20 × 20 tile grid, where each tile counts the average of all spectra integrals stored on the x- versus the y-axis. The inverted ratio is also calculated and mapped to a D3 diverging color mapping function (D3 scale chromatic) for both the heatmap and unit cell detail views. Inside the cell, data are flattened to a single array and integrals of the metabolite regions from x- and y-axes are counted for each set (all data for a voxel, a patient-voxel pair, the patient-voxel-state and finally patient-voxel-state-timepoint) as a measure to avoid duplicate data in calculations.

8. Case Studies

We now present two MRS case studies to provide preliminary evidence for the value and utility of SpectraMosaic in both clinical study and in research settings. Case study data were analyzed by three unpaid, volunteer participants recruited from our group of collaborators from the hospital MR spectroscopy lab group. All three are experienced MR spectroscopy researchers who are not co-authors of this work, and who are unfamiliar with this or similar visualizations. One user is an MR physicist specializing in development and refinement of spectroscopy protocols for clinical studies of neuropsychiatric and developmental disorders. Our second expert uses 31P and 23Na-labeled pyruvate timecourse data to look at real time metabolism. Finally, our third user is interested in using MRS with fMRI to study correlation between these two modalities as well as group differences in neurodegenerative and developmental disorders such as Parkinson's disease, stroke, stuttering, and dyslexia. All three users explored the application, after some brief instruction, with the aim of addressing all tasks outlined in Section 4, following a "think aloud" protocol and follow-on interview after both cases were completed. We then reviewed and collated the feedback from all participants.

Case Study: Difference Assessment Our first case study dataset compares single timepoint spectral acquisitions from a single spatial area at two different echo times (TEs) for a patient cohort. With large changes in echo time it is expected that there will be a major
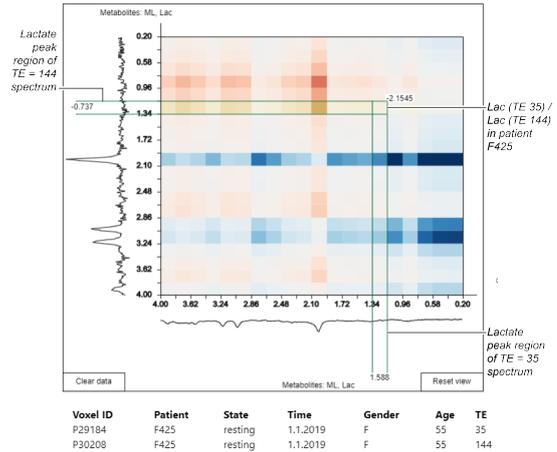

Voxel ID	Patient	State	Time	Gender	Age	TE
P29184	F425	resting	1.1.2019	F	55	35
P30208	F425	resting	1.1.2019	F	55	144

Figure 5: *Detailed inspection in a single-patient, single-voxel sample acquired with 2 different TEs: 35ms (x-axis) and 144ms (y-axis). The user is interested in comparing the ratios of various metabolite peaks, such as lactate:lactate, between the two acquisitions.*

difference in the shape and orientation of spectral peaks – specifically of interest is the shape and orientation of the lactate peak. With this study the user wishes to understand the degree of difference between spectral ratios due to TE changes in the MRS protocol sequence.

In Figure 5 we show a phase of the visual analysis performed by one of our collaborators. Using this case study data they were interested to see how lactate, which typically exhibits a peak inversion at TE=144ms, differs from its shape when captured at a lower TE. Since they were also interested in close inspection of the spectral curves, they chose to compare the different TE spectra for just one patient in order to clearly visualize the line shape. With SpectraMosaic they were able to see the peak inversion of lactate at TE=144ms as well as determine the ratio of lactate at TE=35ms versus lactate at TE=144ms. As their comparison was for a single patient there was no need to expand the heatmap to a more detailed unit cell analysis; this affords an example of the utility of SpectraMosaic at any level of visual analysis.

A major difference case such as this study with different TE acquisitions can mimic a range of neurological and neuropsychiatric pathologies where an entire spectral peak may be absent, a new peak may be introduced, or the shape and orientation of existing peaks may be altered. For example, in brain lesion mapping it is often of interest to understand how biochemical composition of the lesion versus the collateral healthy hemisphere. In Figure 5, the one-patient, two voxel representation makes it easy to visualize which compounds, e.g., Cho/Cr ratio, are elevated and also how they behave at repeated visits. Alternative useful comparisons include using a group average of normal data, or a group average of the expected pathology, to visualize how the single brain lesion sample compares. This could provide insight into the underlying numeric data and provide assistance in staging the disease or monitoring effectiveness of treatment.

Case Study: Similarity Assessment We now present a neuroin-flammation research case study for assessing biochemical similar-ity in multiple spatial voxel locations in multiple patients. With these data the user hopes to answer two primary questions: At a single point in time for a uniform resting state, how similar are con-centrations of various spectral metabolites between spatial regions of the brain, and how similar are they between the two individuals?

Figure 6 demonstrates the analysis process of one user who was interested in comparing spatial variation across two individuals in the average and individual concentration ratio of NAA versus lipids (ML). In this exploration they were able to quickly discern the sub-tle value differences for not only all three voxel locations but also for each patient within these areas of the brain. Although these dif-ferences are small, as indicated by the correspondingly close hue values, one can quickly see that the prefrontal sample in the female patient exhibits the largest ratio value of NAA to lipids at 4.6, while the male patient at that same spatial sample has a ratio of approxi-mately 3.5. Insights such as this led our user to consider the effect that gender, or age, may have in this case, offering new avenues of further exploration for this dataset.

This research study dataset can be considered as a model for a larger class of cases where visual detection of subtle differences are key in the analysis process. For example, in psychiatric research metabolites are mapped to understand how underlying biochemi-cal conditions could influence clinical symptoms. Neurotransmitter concentrations, i.e., glutathione (GSH), glutamate and GABA, can easily vary across individuals, brain states and in response to treat-ment. The subtle changes in concentrations of these metabolites may be investigated on shorter time scales (response to a cognitive task or experimental manipulation such as tDCS) and longer time scales (across imaging sessions). SpectraMosaic can contribute in the assessment of individual variability, variability across time and space (voxel placement) and in understanding the complexity in the data beyond inference on statistical mean values. This constitutes a novel approach to investigate MRS data.

User Feedback All three case study participants expressed en-thusiasm for a novel MRS visual analysis tool and expressed in-terest for its use in their research areas. Our experts felt that the tool was overall intuitive and easy to use – all three liked the drag and drop feature for loading in data and selection of spectral vox-els for heatmap ratio comparisons. One user specializing in clini-cal research expressed interest in seeing clearly in the anatomical image spatial voxel placements layered onto a single normalized brain image as a means to understand sample location consistency over a study cohort, or to enable simultaneous views of all pa-tients for each spatial sample. Conversely, more general research-oriented users preferred the individualized anatomical images with one voxel placed per image. One suggestion was to keep the single image but allow an option for the user to select the voxel image plane: sagittal, axial, or coronal.

Regarding the the spectral heatmap and detail cell visualizations for each metabolite, one user commented, *"It automatically lines up, I do not have to search for it or find some way to align them, next to or on top of each other [when comparing metabolite ra-tios]...this is useful for showing differences clearly and visually.".* Symmetric adjustment of spectral grid tiles was also suggested as
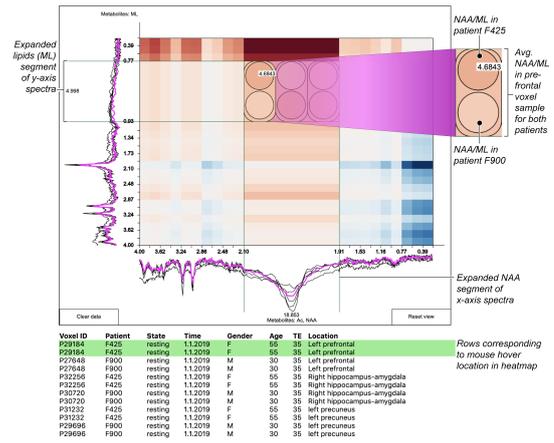

Figure 6: *Similarity case study example where the user has chosen to inspect individual ratio variation for NAA versus lipids (ML) for two patients with three spatial voxels acquired for each patient. In this detail image we can see subtle variation in metabolite concen-tration ratios in each patient, with the greatest difference arising from the prefrontal voxel sample resulting from a higher ratio for patient F425.*

an interaction alternative or addition to the current asymmetric grid adjustment scheme for greater usability. In the visualization of large cohort spectroscopy datasets, one collaborator suggested we apply the individual patient disk visual idiom to also represent patient groupings, serving as a middle layer between our existing first and second visualization tiers. We could then add as option for users to create these groups on-the-fly, which can be flexibly recomposed with changing analysis needs. Our collaborators saw the metabo-lite listings along each axis as a high-value feature: *"In itself it is actually very useful to be able to see where the [metabolite] peak should be, and to see the multiple peaks for metabolites, such as the double peak of creatine.".* The ability to make modifications to this listing set based on analytical goals would help users pinpoint more quickly metabolite(s) of interest for their study. Finally, each user indicated interest for an option to extraction spectral heatmap data as a CSV to use for subsequent statistical analysis. One collabora-tor expressed interest in seeing this output to the hospital PACS for access by radiologists to aid in more rapid interpretation of spec-troscopy data for more widespread use in a clinical setting.

In considering the utility of SpectraMosaic when applied to their specific areas of research, whether clinical research or spectral pro-tocol development, all three of our collaborators felt this could aug-ment their current workflow and provide deeper and more rapid in-sights to their data: *" ...this [spectral heatmap] feature is useful to have a closer look at, for example, neurodegeneration [in Parkin-son's] with the loss of dopaminergic connections, as seen with con-centrations of glutamate or GABA... and it is ideal for testing new protocols against established protocols."* In particular, all felt this tool could be helpful in group comparison, especially in one collab-orator's research in Parkinson's disease. He felt that this tool would be ideal for him to explore and display results of his cohort studies,

where he could then select individuals on which to run subsequent statistical analyses.

9. Discussion

SpectraMosaic currently supports single single or adjacent peak comparison – a logical next step would be supporting comparison of one-to-many (or vice versa) non-contiguous peaks in a spectrum, as these relationships could potentially open up other new insights for our collaborators. This would also address the challenge of metabolites consisting of multiple peaks in a spectrum, e.g., the doublet peak of lactate. As the tool is at the moment, our collaborators have confirmed that it provides a valid, clear and easily read approximation of the ratios and variability between samples for further inquiry. Our tool is also not currently able to compare multiple detail unit cells at once, this would be a helpful analysis feature. Similarly, we also discussed with our collaborators the utility of comparative analysis of two metabolite concentration ratios, i.e., GABA/NAA versus Glx/NAA, within a cell. This is theoretically possible with our visual encoding system where we could recycle the half-moon glyph visualization for state to instead compare two metabolites, but would require reworking the basic setup of the spectral axes to group selected metabolites into unit cells by user selection rather than by chemical shift position.

SpectraMosaic also does not reveal detail glyph visual encoding elements while in the cell overview – this would be an interesting feature extension to allow users to even more rapidly see individual data outliers rather than stepping through to a deeper detail view for these visualizations. Support for uncertainty visualization is also not yet in place as we instead focused our efforts with this tool on developing a consistent core visual hierarchy language for each relevant spectral data dimension. On the topic of visual hierarchy we found after our case studies that TE variation is frequently performed in research for protocol definition and refinement, and classifying this dimension as part of our second, rather than third, visualization tier would be a more appropriate choice. Although our diverging color mapping system is effective in demonstrating large differences, subtle changes important in certain pathological conditions are less obvious and investigation into more fine grained color mapping options or automatic or user defined color map scaling may help more clearly highlight these micro-changes in tissue metabolic concentrations.

10. Conclusions and Future Work

In this design study we contributed a thorough characterization of the data, task, and design requirements inherent to a successful spectroscopy visualization tool. We followed this with presentation of our design rationale and visual encodings for SpectraMosaic, a novel visualization approach for hypothesis-driven MRS data comparison and exploration. Finally, we performed case studies with three domain experts to validate our tool in spectroscopy clinical and research studies. MRS is a ripe area for continued visualization research, and we see a number of opportunities for expansion and refinement of SpectraMosaic.

Although typical MRS studies are small (< 15 patients), a logical

extension is the expansion of our visual encoding system to successfully manage larger cohorts. We envision this is a possibility with the creation of analysis groups that would form a visualization tier between our previously established first and second tiers. It would also be interesting for future extensions of this tool to permit calculations and comparisons of more complex metabolite relationships, both outside and inside the unit cell visualizations. Addition of a visual idiom for data uncertainty characterization could provide further validation of acquisition quality. Lastly, the flexible design of our tool is such that new statistical measures, e.g., correlation coefficient, may be swapped in or added to the space, affording further insights into the data. We furthermore plan to extend this application for analysis of other elements used in clinical spectroscopy research, including 31P and 23Na – this can be easily done with the extension of our CPK color idiom with orange and blue representing phosphorus and sodium, respectively.

We plan to continue working with our collaborators to further refine SpectraMosaic and have deployed this tool *(https://folk.uib.no/lga066/spectramosaic/)* for their use. Source code for the application is also available for pull requests *(https://git.app.uib.no/Laura.Garrison/spectramosaic)*. Beyond the medical space, an additional interesting line of inquiry would be in exploring the adaptability of our abstracted tasks paired with our visual encoding system in other domains facing similar challenges with heterogeneous multidimensional data, such as meteorology or geology.

References

[AFK*14] AMIRKHANOV A., FRÖHLER B., KASTNER J., GRÖLLER E., HEINZL C.: InSpectr: Multi-modal exploration, visualization, and analysis of spectral data. *Computer Graphics Forum 33*, 3 (2014), 91–100. doi:10.1111/cgf.12365. 2

[BBW83] BERTIN J., BERG W. J., WAINER H.: *Semiology of graphics: diagrams, networks, maps*. University of Wisconsin press Madison, 1983. doi:10.1080/00690805.1987.10438353. 2

[BM10] BRUCKNER S., MÖLLER T.: Isosurface similarity maps. *Computer Graphics Forum 29*, 3 (2010), 773–782. doi:10.1111/j.1467-8659.2009.01689.x. 2

[BM13] BREHMER M., MUNZNER T.: A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2376–2385. doi:10.1109/TVCG.2013.124. 3

[CON13] CRANE J. C., OLSON M. P., NELSON S. J.: SIVIC: open-source, standards-based software for DICOM MR spectroscopy workflows. *Journal of Biomedical Imaging 2013* (2013), 12. doi:10.1155/2013/169526. 3

[FKLT10a] FENG D., KWOCK L., LEE Y., TAYLOR R.: Matching visual saliency to confidence in plots of uncertain data. *IEEE Transactions*

on Visualization and Computer Graphics 16, 6 (2010), 980–989. `doi:10.1109/TVCG.2010.176`. 2

[FKLT10b] FENG D., KWOCK L., LEE Y., TAYLOR R. M.: Linked exploratory visualizations for uncertain MR spectroscopy data. *Visualization and Data Analysis 7530* (2010), 753004. `doi:10.1117/12.839818`. 2

[FLKT09] FENG D., LEE Y., KWOCK L., TAYLOR R. M.: Evaluation of glyph-based multivariate scalar volume visualization techniques. In *Proceedings of the Symposium on Applied Perception in Graphics and Visualization* (2009), pp. 61–68. `doi:10.1145/1620993.1621006`. 2

[GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Information Visualization 10*, 4 (2011), 289–309. `doi:10.1177/1473871611416549`. 2

[HB03] HARROWER M., BREWER C. A.: ColorBrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal 40*, 1 (2003), 27–37. `doi:10.4324/9781351191234-18`. 5

[Her06] HERRAEZ A.: Biomolecules in the computer: Jmol to the rescue. *Biochemistry and Molecular Biology Education 34*, 4 (2006), 255–261. `doi:10.1002/bmb.2006.494034042644`. 6

[LSBP18] LAWONN K., SMIT N. N., BÜHLER K., PREIM B.: A survey on multimodal medical data visualization. *Computer Graphics Forum 37*, 1 (2018), 413–438. `doi:10.1111/cgf.13306`. 3

[MFGH08] MATKOVIC K., FREILER W., GRACANIN D., HAUSER H.: ComVis: A coordinated multiple views system for prototyping new visualization technology. In *International Conference Information Visualisation* (2008), pp. 215–220. `doi:10.1109/IV.2008.87`. 2

[MK11] MARINO J., KAUFMAN A.: Prostate cancer visualization from MR imagery and MR spectroscopy. *Computer Graphics Forum 30*, 3 (2011), 1051–1060. `doi:10.1111/j.1467-8659.2011.01954.x`. 2

[MMDP10] MEYER M., MUNZNER T., DEPACE A., PFISTER H.: MulteeSum: a tool for comparative spatial and temporal gene expression data. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 908–917. `doi:10.1109/TVCG.2010.137`. 2

[MMO*14] MULLINS P. G., MCGONIGLE D. J., O'GORMAN R. L., PUTS N. A., VIDYASAGAR R., EVANS C. J., EDDEN R. A. E., ET AL.: Current practice in the use of MEGA-PRESS spectroscopy for the detection of GABA. *Neuroimage 86*, 1 (2014), 43–52. `doi:10.1016/j.neuroimage.2012.12.004`. 3

[MWS*10] MEYER M., WONG B., STYCZYNSKI M., MUNZNER T., PFISTER H.: Pathline: A tool for comparative functional genomics. *Computer Graphics Forum 29*, 3 (2010), 1043–1052. `doi:10.1111/j.1467-8659.2009.01710.x`. 2, 5

[NLK*14] NUNES M., LARUELO A., KEN S., LAPRIE A., BÜHLER K.: A survey on visualizing magnetic resonance spectroscopy data. In *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine* (2014), pp. 21–30. `doi:10.2312/vcbm.20141180`. 2

[NRS*14] NUNES M., ROWLAND B., SCHLACHTER M., KEN S., MATKOVIC K., LAPRIE A., BÜHLER K.: An integrated visual analysis system for fusing MR spectroscopy and multi-modal radiology imaging. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), pp. 53–62. `doi:10.1109/VAST.2014.7042481`. 2

[PCB*17] PURVIS L. A., CLARKE W. T., BIASIOLLI L., VALKOVIČ L., ROBSON M. D., RODGERS C. T.: OXSA: An open-source magnetic resonance spectroscopy analysis toolbox in MATLAB. *PloS one 12*, 9 (2017), e0185356. `doi:10.1371/journal.pone.0185356`. 3

[PDF14] PERIN C., DRAGICEVIC P., FEKETE J.: Bertifier: New interactions for crafting tabular visualizations. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2082 – 2091. `doi:10.1109/TVCG.2014.2346279`. 6

[PDFE17] PARK D., DRUCKER S., FERNANDEZ R., ELMQVIST N.: Atom: A grammar for unit visualizations. *IEEE Transactions on Visualization and Computer Graphics 24*, 12 (2017), 3032 – 3043. `doi:10.1109/TVCG.2017.2785807`. 2

[PODA13] POSSE S., OTAZO R., DAGER S. R., ALGER J.: MR spectroscopic imaging: Principles and recent advances. *Journal of Magnetic Resonance Imaging 37* (2013), 1301–1325. `doi:10.1002/jmri.23945`. 5

[Pro01] PROVENCHER S. W.: Automatic quantitation of localized in vivo 1H spectra with LCModel. *NMR in Biomedicine 14*, 4 (2001), 260–264. `doi:10.1002/nbm.698`. 3

[SDCA*09] STEFAN D., DI CESARE F., ANDRASESCU A., POPA E., LAZARIEV A., VESCOVO E., STRBAK O., WILLIAMS S., STARCUK Z., CABANAS M., ET AL.: Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package. *Measurement Science and Technology 20*, 10 (2009), 104035. `doi:10.1088/0957-0233/20/10/104035`. 3

[SMW95] SAYLE R. A., MILNER-WHITE E. J.: RASMOL: biomolecular graphics for all. *Trends in biochemical sciences 20*, 9 (1995), 374–376. `doi:10.1016/S0968-0004(00)89080-5`. 6

[SR13] STAGG C., ROTHMAN D. L.: *Magnetic resonance spectroscopy: tools for neuroscience research and emerging clinical applications.* Academic Press, 2013. 3, 4, 5

[STH02] STOLTE C., TANG D., HANRAHAN P.: Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Transactions on Visualization and Computer Graphics 8*, 1 (2002), 52–65. `doi:10.1145/1400214.1400234`. 2

[UBA16] ULMER S., BACKENS M., AHLHELM F. J.: Basic principles and clinical applications of magnetic resonance spectroscopy in neuroradiology. *Journal of Computer Assisted Tomography 40*, 1 (2016), 1–13. `doi:10.1097/RCT.0000000000000322`. 3

[VDG10] VAN DER GRAAF M.: In vivo magnetic resonance spectroscopy: Basic methodology and clinical applications. *European Biophysics Journal 39*, 4 (2010), 527–540. `doi:10.1002/9780470882221`. 1, 3

[VKV*16] VERMA A., KUMAR I., VERMA N., AGGARWAL P., OJHA R.: Magnetic resonance spectroscopy - revisiting the biochemical and molecular milieu of brain tumors. *BBA Clinical 5* (2016), 170–178. `doi:10.1016/j.bbacli.2016.04.002`. 2

[War94] WARD M. O.: Xmdvtool: integrating multiple methods for visualizing multivariate data. In *Proceedings of the Conference on Visualization* (1994), IEEE Computer Society Press, pp. 326–333. `doi:10.1109/visual.1994.346302`. 5

[WLT94] WARD M., LEBLANC J. T., TIPNIS R.: N-land: a graphical tool for exploring n-dimensional data. In *Proceedings of Computer Graphics International Conference* (1994), pp. 131–141. 5

[WRK*11] WILSON M., REYNOLDS G., KAUPPINEN R. A., ARVANITIS T. N., PEET A.: A constrained least-squares approach to the automated quantitation of in vivo 1H magnetic resonance spectroscopy data. *Magnetic resonance in medicine 65*, 1 (2011), 1–12. `doi:10.1002/mrm.22579`. 3

[WVW*05] WOLF I., VETTER M., WEGNER I., BÖTTGER T., NOLDEN M., SCHÖBINGER M., HASTENTEUFEL M., KUNERT T., MEINZER H. P.: The medical imaging interaction toolkit. *Medical Image Analysis 9*, 6 (2005), 594–604. `doi:10.1016/j.media.2005.04.005`. 2

[XV10] XU D., VIGNERON D.: Magnetic resonance spectroscopy imaging of the newborn brainâĂȚa technical review. *Seminars in perinatology 34*, 1 (2010), 20–27. `doi:10.1053/j.semperi.2009.10.003`. 4