

Special Section on VCBM 2019

# Interactive visual exploration of metabolite ratios in MR spectroscopy studies



Laura Garrison<sup>a,c,\*</sup>, Jakub Vašíček<sup>b</sup>, Alexander R. Craven<sup>e,f</sup>, Renate Grüner<sup>c,d</sup>,  
Noeska N. Smit<sup>a,c</sup>, Stefan Bruckner<sup>a,c</sup>

<sup>a</sup> Department of Informatics, University of Bergen, Thormøhlens Gate 55, Bergen 5008, Norway

<sup>b</sup> Department of Informatics, Masaryk University, Brno, Czech Republic

<sup>c</sup> Department of Radiology, Mohn Medical Imaging and Visualization Centre, Haukeland University of Hospital, Bergen, Norway

<sup>d</sup> Department of Physics and Technology, University of Bergen, Norway

<sup>e</sup> Department of Biological and Medical Psychology, University of Bergen, Norway

<sup>f</sup> Department of Clinical Engineering, Haukeland University of Hospital, Bergen, Norway

## ARTICLE INFO

### Article history:

Received 17 January 2020

Revised 29 June 2020

Accepted 6 August 2020

Available online 13 August 2020

### Keywords:

Computing methodologies and applications

Information systems (hypertext navigation

interfaces

decision support

etc.)

Applications to biology and medical

sciences

Biomedical imaging and signal processing

## ABSTRACT

Magnetic resonance spectroscopy (MRS) is an advanced biochemical technique used to identify metabolic compounds in living tissue. While its sensitivity and specificity to chemical imbalances render it a valuable tool in clinical assessment, the results from this modality are abstract and difficult to interpret. With this design study we characterized and explored the tasks and requirements for evaluating these data from the perspective of a MRS research specialist. Our resulting tool, SpectraMosaic, links with upstream spectroscopy quantification software to provide a means for precise interactive visual analysis of metabolites with both single- and multi-peak spectral signatures. Using a layered visual approach, SpectraMosaic allows researchers to analyze any permutation of metabolites in ratio form for an entire cohort, or by sample region, individual, acquisition date, or brain activity status at the time of acquisition. A case study with three MRS researchers demonstrates the utility of our approach in rapid and iterative spectral data analysis.

© 2020 The Author(s). Published by Elsevier Ltd.

This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Magnetic resonance spectroscopy (MRS) is an *in vivo* non-invasive biochemical technique used to estimate the concentrations of certain small molecules, known as metabolites, in a tissue region. When paired with high structural resolution MR imaging (MRI), it has shown clinical potential for improving diagnosis and treatment monitoring of numerous diseases and disorders of the central nervous system [1]. However, its clinical adoption remains limited. Translation from the metabolite signals acquired from MRS into clinically useful biomarkers is an open challenge in spectroscopy research. Optimization and tuning of parameters for consistent, isolated metabolite acquisition is one such area of

research, while another branch of research aims to identify patterns of the subtle disease effects on multiple metabolites [2]. In this paper, we explore the application of visualization techniques to identify ratios and patterns of multiple metabolites.

While recent technology improvements in MRS acquisition have enhanced data quality and resolution [3], visualization of MRS data remains a largely unexplored area. MR spectroscopy produces a vastly different readout than MR imaging. Rather than a greyscale image of recognizable anatomical structures over many voxels, it acquires an abstract spectrum per single voxel. This spectrum consists of a series of peaks (resonances) that represent signal intensities as a function of frequency, as depicted in Fig. 1. Metabolites may consist of single peaks, as in the case of N-acetylaspartate (NAA), or multiple peaks, as in Creatine (Cr). Most tools used to quantify single voxel spectral data, e.g., LCModel [4] produce only rudimentary visual output, such as the spectral graph in Fig. 1. Recognizing the metabolites that correspond to these graphs is challenging. Although it is important to see the spectral graph as a means of quality assurance, metabolite concentrations are the most

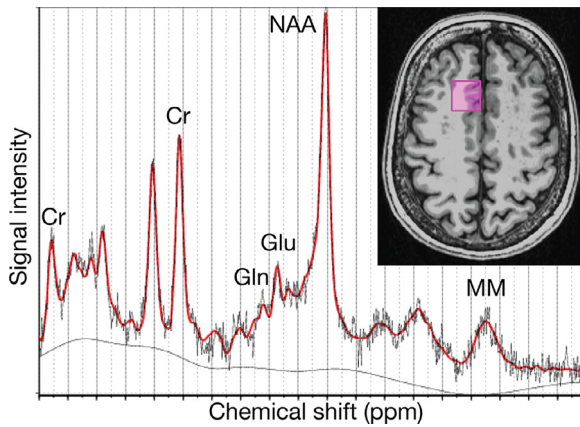
This article has been certified as Replicable by the *Graphics Replicability Stamp Initiative*: <http://www.replicabilitystamp.org>

\* Corresponding author at: Department of Informatics, University of Bergen, Thormøhlens Gate 55, 5008 Bergen, Norway.

E-mail address: [laura.garrison@uib.no](mailto:laura.garrison@uib.no) (L. Garrison).

<https://doi.org/10.1016/j.cag.2020.08.001>

0097-8493/© 2020 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)



**Fig. 1.** Plot of the typical visual output for spectral quantification by LCMoel, with processed data and model (overlaid smooth curve). Structural localization and annotations of key metabolite peaks have been added.

clinically relevant output from this method. These concentrations are most often output to a simple table in standard domain tools. This does little to advance interpretation, understanding, or to facilitate rapid comparison of metabolites between acquisitions.

This paper expands upon our previous design study [5] in building a general tool for the interactive visual analysis of all permutations of spectral metabolites, in ratio form, for a small cohort. While we previously emphasized rapid visualization of metabolite ratios directly from spectral input data, this work allows visualization of complete, and more complex, metabolic signatures via an integrated pipeline with Tarquin [6], an open source spectral quantification tool. Our specific contributions include:

1. We provide a detailed review of MRS data characteristics and abstraction of spectral analysis tasks identified from domain expert collaboration.
2. We present a refined pipeline that integrates spectral quantification and fitting to allow multi-peak metabolite analysis.
3. Our visual exploratory analysis tool provides an extended interface for linking of structural, spectral, and patient data, including group creation and uncertainty communication.
4. We introduce a tiered system of visual encodings depicting layers of aggregated metabolite ratios that can be partitioned by key attributes.
5. We present a clinical case study and feedback from three MR spectroscopy research experts.

Using SpectraMosaic, MR spectroscopy researchers are able to rapidly identify patterns at different layers that may be of interest for deeper clinical exploration.

## 2. Related work

A key challenge in visualizing spectroscopy data is that each spectrum is in itself a multivariate dataset. We draw inspiration from tools such as InSpectr [7], which utilizes multiple linked views and comparative visualization techniques [8] from multimodal data sources (x-ray computed tomography and x-ray fluoroscopy) to provide insights into composition of a multivariate sample. SpectraMosaic similarly combines imaging techniques (MRS and MRI), but for a different domain and with a different focus. Isosurface similarity maps defined by Bruckner and Möller [9] were applied to spectra in Spectral Similarity Maps, an extension of the InSpectr framework [10]. In this approach, 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 metabolite ratios.

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) presented by Feng et al. [11] provide a 3D representation of metabolites within a voxel. This application was later extended to include scatter and parallel coordinate plots for a subset of metabolites [12]. SpectraMosaic remains in the abstract visualization space, but allows comparison of all metabolite ratios. Nunes et al. [13] presented a visual analysis framework combining ComVis [14] and MITK [15]. Brushing and linking mechanisms allow for the definition of a biological target volume with its corresponding metabolite values. However, this work was developed specifically for radiotherapy treatment visualization. Retention of spectra was not the focus of the application and it provided limited functionality for metabolite comparison. SpectraMosaic extends the flexibility of metabolite ratio calculations, and displays additional MRS data attributes (spatial, individual, temporal, and brain activity status) in an overview and detail visual representation. Marino and Kaufman [16] implemented direct volume rendering (DVR) to represent male prostate anatomy from MRI data combined with PET and MRS in prostate tumor delineation. However, this application was focused on a single metabolite ratio, and could only present an individual in a single time slice. SpectraMosaic retains an abstract visualization format, but offers broader insights into metabolite relationships over time and between individuals. Jawad et al. [17] developed a system for the analysis of segmented brain tissue composition to identify the metabolic signatures of brain tumors—this tool was optimized for multivoxel data, and focused on statistical outcome measurements. SpectraMosaic works at a more generalized level in spectral analysis. Further work by Jawad et al. [18] presented an approach for the comparative analysis of single voxel spectroscopy in cohort data, focusing primarily on violin and parallel coordinate plots to convey spectral metabolite relationships. Our approach uses a similar range of data inputs and processing tools. However, our tool focuses on simultaneous comparison of all metabolite ratios, using a nested visual design linking multiple MRS data elements.

First introduced by Bertin [19], numerous solutions have leveraged small related graphics series to visualize multivariate data. We base SpectraMosaic on this concept, but extend this by including a second layer of nested visual encodings. This is inspired by AtOM [20], a grammar for unit visualizations where individual data items are represented by unique visual marks (units) in a visual encoding system. PivotTable, subsequently trademarked by Microsoft and extended by Polaris [21], enables exploration and analysis of multidimensional data with the flexibility to modify visual encodings, graphics, and table configuration for visualization. Klemm et al. [22] built on this concept for linked visualization of image-centric heterogeneous cohort data. Our approach is related in that we allow on-the-fly reconfiguration of our matrix inputs. Although the cohorts our application focuses on are not large, we share similar considerations with heterogeneous and multivariate data inputs.

While our prior iteration of the SpectraMosaic application focused on the rapid analysis of single-peak metabolites directly from spectral graphs [5], this work expands the tool to allow full, precise spectral analysis in an integrated pipeline with robust MR spectroscopy quantification tools. This permits analysis of metabolites with more complex metabolic signatures; these are encoded to bar and box plots for ease of interpretation. We further increase the practical usability of the tool with new facilities for analysis group creation and additional means for conveying the underlying data distribution. These features arose from additional working sessions and discussions with spectroscopy researchers.

### 3. Background

MRS is an advanced spectroscopic technique used to non-invasively describe the biochemical composition of living tissue. While MRI shows the spatial distribution of atomic nuclei with high spatial resolution, MRS trades spatial resolution for detailed chemical information, using the same hardware. For example, where MRI may be used to identify the extent of a tumor, MRS can help to identify the type of tumor [23]. For each measured voxel, MRS produces a spectrum of signal intensity as a function of frequency. Intensity peaks at different resonance frequencies are described as chemical shifts. These chemical shifts, expressed in parts per million (ppm), arise from fundamentally different nuclear properties of the chemical structures being measured, and represent metabolites in the acquired voxel [24]. The most commonly measured signal comes from hydrogen atoms; this is known as proton MRS ( $^1\text{H-MRS}$ ). This technique is capable of detecting metabolites in concentrations 50,000 times lower than that of fat or water as imaged in conventional MRI.

MRS acquisition techniques include single voxel spectroscopy (SVS) or chemical shift imaging (CSI). CSI is essentially a slab of multiple smaller single voxels. It covers a much larger spatial area than SVS, but suffers from a reduced signal-to-noise ratio. CSI produces a low-resolution image for each metabolite, being in that way similar to conventional MRI, while SVS is more abstract and cannot be visualised in a conventional way. Since SVS acquisition techniques afford more detailed spectra for analysis, we focus our work on this technique. The majority of acquisitions by our collaborators are collected at single time points, i.e., in longitudinal studies, but may also be captured as time-resolved concentrations within a single examination, i.e., functional studies. In the latter approach the subject can also be asked to perform tasks, such as tapping fingers during the acquisition (active brain state), and alternately resting (resting brain state).

Following acquisition, data are output to a vendor-specific format that contains raw data and a header file containing all experimental parameters. Subsequent preprocessing and quantification steps follow to map spectral peak intensities to metabolite concentrations in the measured voxel. In a final fitting step, a model based on prior information is fit to the acquired spectrum; in many approaches, this is effectively a linear combination of basis sets consisting of simulated or measured metabolite signatures. Metabolite concentrations are typically calculated relative to a stable reference, often water or creatine. This allows for a direct comparison of relative metabolite concentrations, assuming the same acquisition hardware and protocols are used. While a more comprehensive discussion of all steps is beyond the scope of this paper, interested readers can refer to Stagg et al. [25] for a detailed overview. A number of existing tools can be used to perform these steps: LCModel [4] is one such widely-used commercial tool, while jMRUI [26], TARQUIN [27], SIVIC [28], OXSA [29], and Gannet [30] offer open source solutions. Equipment manufacturers also supply basic tools to facilitate simple analyses on the scanner console. Our collaborators typically use LCModel or Tarquin; we utilize Tarquin in our pipeline for its ease of use and open availability. The output from these steps includes the experimental parameters used for the acquisition as well as the fitted data and quantification information for each metabolite.

### 4. Task and requirement analysis

We developed SpectraMosaic over the course of one year. We met weekly with our domain collaborators, two of whom are coauthors of this paper. Collaborator backgrounds included two MD/PhDs in radiology, eight PhD researchers in MR imaging, and three MR engineers. The weekly meetings went through three dis-

tinct phases. The first phase focused on domain evaluation, identification of key challenges and where visualisation could potentially help overcome them. Ultimately, the output from this phase was agreement on core tasks and requirements. The second phase explored the design space for these tasks/requirements with discussion and interface prototypes. These were refined and narrowed down to a single option. Our third phase reviewed and refined an alpha application. Basic use case testing alongside individual and group evaluation feedback ultimately helped us settle on the version we present in this paper.

#### 4.1. Task analysis and abstraction

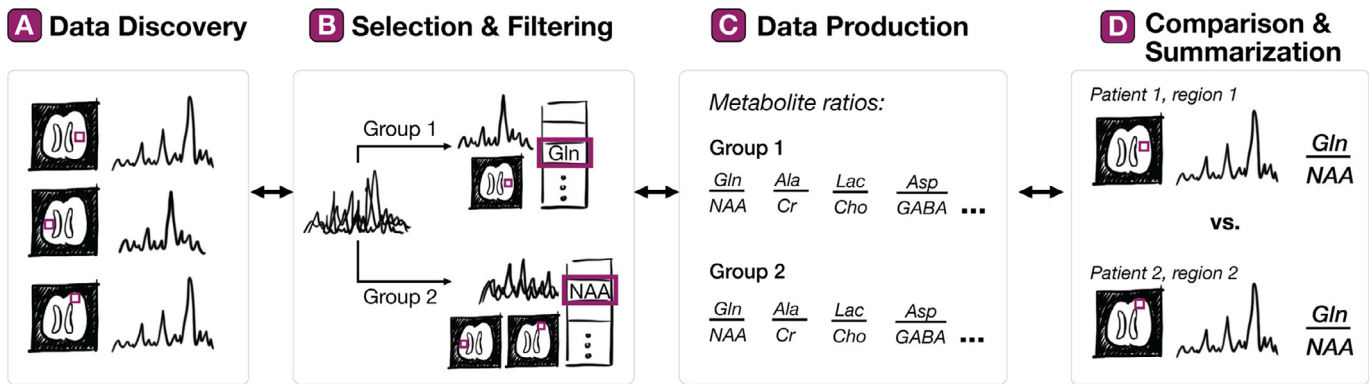
We frame the analysis tasks identified in phase one of our collaboration in the context of Brehmer and Munzner's multi-level typology of abstract visualization tasks [31]. This abstraction was useful for our development process, as it allowed us to more objectively frame the challenges experienced by our colleagues. These tasks form a generalized workflow shown in Fig. 2. The first step, data discovery, provides a general overview of the input components for spectral analysis. Following user selection of components for analysis, a data production step calculates ratios from all inputs. Ratio comparison and summarization follows.

*T1: Data discovery.* The first set of tasks relates to data consumption for discovery and verification of key MRS data aspects (Fig. 2(A)). Spectra, anatomical reference images, and associated subject data are reviewed together in an initial overview step. Researchers visualize spectral graphs to establish a general sense of the data quality and to form initial hypotheses. Supplemental parameter information, such as the echo time (TE), during the acquisition can be used to verify validity of experimental comparisons. Researchers additionally validate their assumptions about the spectral graph against its sample location. This serves two purposes: (1) as a second quality assurance measure to check whether the data were sampled in the correct region, and (2) to provide initial validation for graph differences between spatial regions. This is because a normal spectrum in one area of the brain may be aberrant in another region with a different tissue composition [25].

*T2: Selection and filtering.* Following an overview, researchers next select and filter the data (Fig. 2(B)). In both medical and clinical research studies our collaborators often wish to select a subset of spectra or metabolites for further analysis for a variety of different reasons. For instance, researchers may wish to look only at the variation in metabolite concentration ratios for a single time acquisition in a longitudinal cohort study, e.g., pre-operative patients in a tumor cohort, or to analyze only female subjects within a study. Furthermore, some metabolites may be uninteresting to include for certain clinical studies, e.g., lipids and macromolecules are not usually relevant outside of certain oncological studies [25], and are useful to exclude on-demand.

*T3: Data production.* Spectra can vary considerably between acquisitions. This can occur due to different acquisition parameter settings or simply between different scanners. Ratios and correlations calculated from metabolite concentrations are two standard methods to understand spectroscopy data [32]. The use of ratios to determine metabolite concentrations is a core critical task for any MRS application for two reasons, (1) as a method to correct for inhomogeneity across the sample and (2) to account for varying tissue composition. Following selection of interesting metabolites for analysis, a data derivation step takes as input the Tarquin-processed and quantified metabolite values and outputs the metabolites in ratio form (Fig. 2(C)).

*T4: Comparison and summarization.* Following data derivation, researchers then wish to summarize and compare metabolite ratios (Fig. 2(D)). For example, researchers studying oxygen deprivation (hypoxia) in newborns are interested in comparing the metabolic



**Fig. 2.** Typical task flow for MRS data analysis. Users begin with data discovery (A) to review spatial voxel position, associated spectral graphs, and relevant acquisition parameters. (B) continues with data selection and filtering, where spectral voxels of interest are selected and divided into groups. Data production (C) calculates all possible ratios of selected metabolites, e.g., Glutamine (Gln) to N-acetylaspartate (NAA). In (D) ratios are compared and summarized between, e.g., Gln/NAA for different patients or different brain regions. Each of these steps may be revisited.

differences between healthy and hypoxic newborns. This can be achieved by evaluating ratios of the same metabolites between both groups, e.g., NAA healthy vs. NAA hypoxic. Furthermore, researchers would like to understand the metabolic profile of hypoxic newborns on a spatial and individual level. For instance, the basal ganglia region of the brain is known to be sensitive to oxygen deprivation, so it is clinically relevant to compare this region to a less sensitive region. Within a given region of interest researchers then wish to compare individuals to identify clinically relevant outliers in order to answer questions such as “How does Lactate/Choline compare for Patient X versus Y?” Moreover, oxygen-deprived newborns who survive often experience developmental disabilities later in life. Longitudinal MRS studies allow researchers to understand how the metabolic profiles of affected individuals change over time relative to healthy individuals. In a different scenario, researchers studying schizophrenia are interested in comparing the metabolic profiles of individuals when their brains are *active* relative to their *resting* brain state. Different metabolites present in different concentrations in these states, and identification of these differences may help progress understanding of this disorder.

Following comparison of interesting metabolite ratios, researchers often wish to refine their hypotheses and revisit metabolic input data. This task sequence then repeats, following an iterative analytical approach to hypothesis exploration and verification in MRS data.

#### 4.2. Design requirements

Following the identification of tasks important for our collaborators in MRS analysis, we developed the design requirements for our application. First, on a technical and infrastructure level, our colleagues often switch between hospital workstations while accessing sensitive patient data. Thus, for practical utility it is critical to provide a tool that enables a *machine-independent workflow* (R1) that adheres to *patient data restrictions* (R2).

As discussed in T3, for a combined analysis of spectra acquired from different scanners, or with different acquisition settings, it is necessary to calculate *metabolite ratios* (R3). Furthermore, as implied by T1 and T2, *visual linking between input data* (voxel placement, spectral graph, patient- and acquisition-specific information) and *calculated metabolite ratios* is important for many analysis questions (R4). For our collaborators, the most important patient and scanner-specific information to retain include patient age, gender, and echo time (TE).

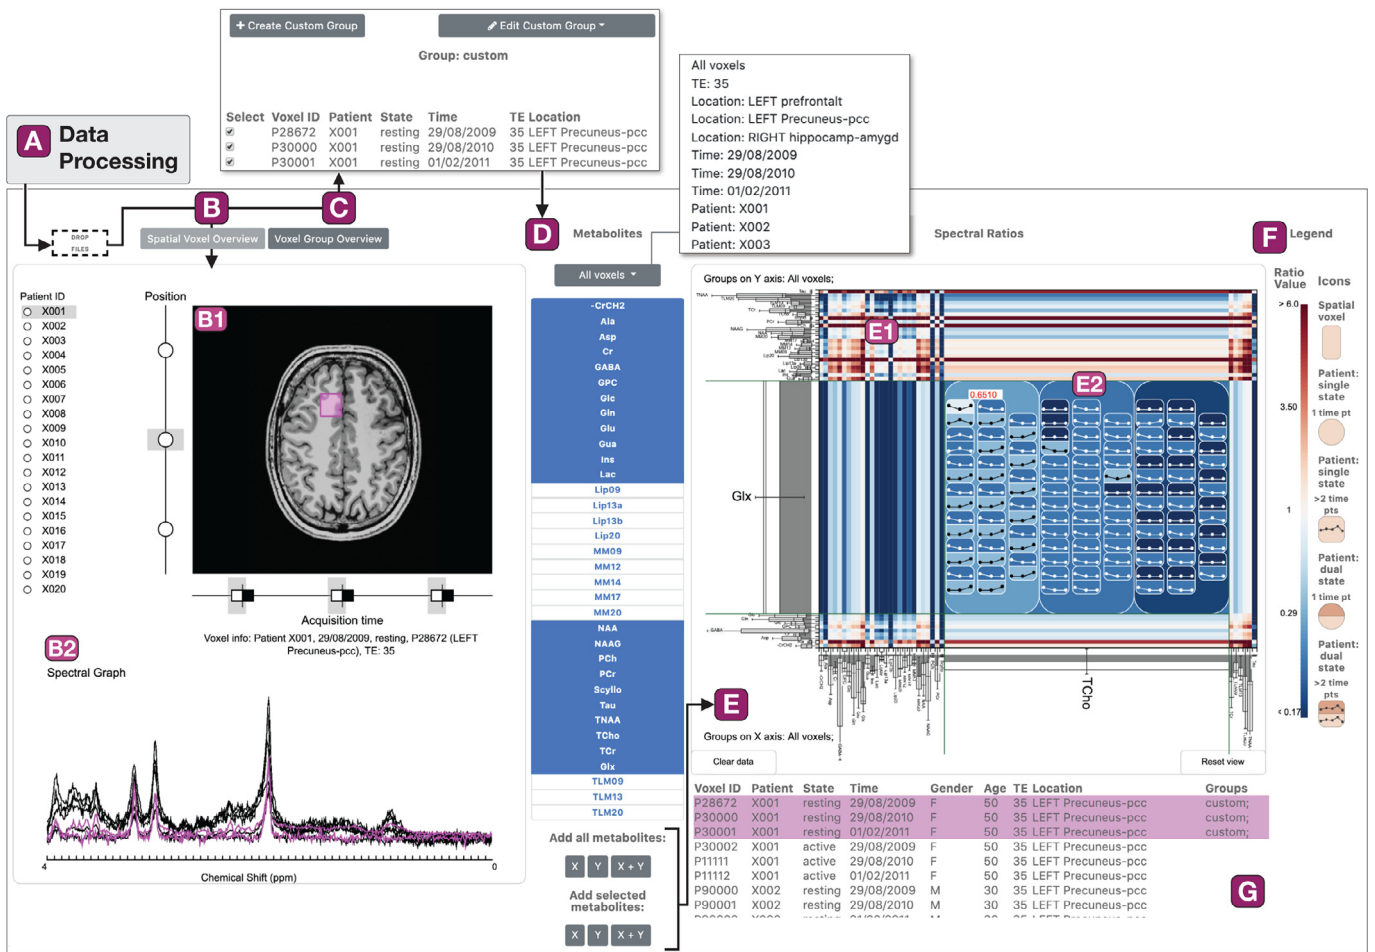
Based on the types of questions outlined in T4, users must be able to *compare metabolite ratios of interest* (R5). This should be accomplished for any permutation through four key attributes: *spatial region, individual, time point, and brain activity state*. Additionally, appropriate mechanisms to compare ratios over time as well as between spatial regions and individuals are critical for longitudinal or single-run studies. Furthermore, for functional MRS studies it is important to support comparison of metabolite ratios in an active relative to a resting brain state.

#### 5. SpectraMosaic workflow and interface

We provide an overview of the SpectraMosaic interface in Fig. 3. Following an offline processing step, data are loaded into the web tool (Fig. 3(A)). Data of interest for analysis can be explored, selected, and added (Fig. 3(B)–(D)) to a spectral ratio heatmap for deeper inquiry and hypothesis verification (Fig. 3(E)). A legend provides information on the encodings used in the tool (Fig. 3(F)). A table below the heatmap summarizes salient acquisition information (Fig. 3(G)).

*Data processing and loading.* We first perform an offline processing step that automates spectral processing and quantification from Tarquin and MATLAB [33]. We utilize MATLAB to process the structural imaging files, which includes patient data de-identification (R2). The resulting output contains a structural image to localize the voxel sample, the spectral graph, quantified metabolites, and associated metadata; these data remain semantically linked in the visual tool. We use a custom data format because the DICOM standard is not universally or consistently adopted for MRS data.

*Visual inspection of voxel positioning and spectral graphs.* Following data loading (Fig. 3(A)), the spatial voxel overview panel (Fig. 3(B)) is used to review the spectral graph, associated anatomical image, and included metadata for each acquisition. This panel consists of a set of images for each patient. In each structural image, a fuchsia rectangle indicates the voxel sample region for the MRS acquisition (Fig. 3(B1)). To the left, a position selector consists of small filled nodes, each of which indicates an acquisition for the selected patient. Using the standard CPK color convention for atomic elements [34], we represent  $^1\text{H}$  spectral metabolites with a white-filled node. A light gray bar behind the disks shows the active selection image, while the node becomes filled in fuchsia to indicate image linkage to a spectrum that is selected in the spectral heatmap panel (Fig. 3(B2)). Users can access different images via these position nodes or time acquisition nodes (horizontal



**Fig. 3.** SpectraMosaic application workflow overview. Raw spectral data are first processed in an offline step (A), then loaded into the application. In (B) the user visualizes the anatomical image with voxel placement for each acquisition (B1) and the associated spectral graph (B2). In (C) users may create custom groups for analysis. Metabolites may be selected (D) for analysis from custom or preset groups in a drop-down list, and selections assigned to the x- or y-axis of a ratio heatmap (E). The ratio heatmap is divided into a cell grid (E1) based on the number of metabolite inputs to each axis. Detailed inspection of a cell (E2) shows the ratios in a series of nested glyphs representing spatial region, individual, individual brain state, and individual time acquisitions. A legend at the right provides a reference for heatmap glyphs and colors (F). A table (G) shows relevant metadata for each voxel.

axis). A selected node shows the structural image with localized voxel, associated spectral graph, and supplemental metadata, such as TE setting, patient age and gender, stored with that voxel (R4). These data are stored hierarchically, where each voxel sample with spatial information is first sorted by individual identifier and associated metadata, then by time of acquisition, and finally by brain activity state during the acquisition.


**Group creation and metabolite selection.** Following visual inspection of voxel position and spectral graphs, the user may then create custom groups of spectral voxels for subsequent analysis (T3) in the Voxel Group Overview panel (Fig. 3(C)). Custom groups may be edited at any time. Membership in a custom group is listed in the metadata table at the bottom right region of the interface (Fig. 3(G)). Our application additionally creates preset groups for each echo time, spatial region, individual, brain state, and time point. These may be immediately accessed in a drop-down list in the Metabolite panel (Fig. 3(D)).

Following a group selection from the metabolite drop-down list, all quantified metabolites from the offline processing step are displayed. Users then have the option of adding all metabolites in the list to the x-axis, y-axis or both axes of a spectral ratio heatmap located to the right of this list (Fig. 3(E)). Alternatively, only a subset of metabolites may be added to the heatmap axes. Groups may be

flexibly added or removed from either axis at any time. Metabolites populate along heatmap axes in alphabetical order; we discussed a number of ordering options with our domain collaborators, settling on this ordering method for consistency and pattern recognition between studies.

**Ratio exploration.** Following loading of metabolite groups onto each axis, we determine ratios for all metabolite permutations for display in the heatmap panel (Fig. 3(E)). This serves as the primary visualization component of our tool, as shown in Fig. 3(E) and which is described in detail in Section 6. In this view, users can compare average (Fig. 3(E1)) or individual metabolite ratios at different levels of detail (R5). Users may interactively expand a cell to reveal key attribute details (Fig. 3(E2)), as inspired by Bertifier [35]. The background of the cell remains visible behind individual ratio elements for all expansions to preserve context of the aggregated value during navigation. This subtle context preservation was deemed useful by experts in our development process.

A legend at the far right (Fig. 3(F)) serves to indicate hue and glyph meaning. Hovering over a cell or glyph correspondingly highlights linked data elements in fuchsia, including the associated spectral graph, patient anatomical image, and associated metadata (R4), as depicted in Fig. 3G.



**Fig. 4.** Box and bar plots encode metabolite concentrations (tier 1 visual encoding). (A) utilizes bars where height encodes the concentration of each metabolite for a single spectral input. (B) is utilized for two to four spectral inputs on an axis, where height encodes the median value and whiskers encode the minimum and maximum metabolite concentration values, respectively. A box plot (C) is utilized for five or more spectral inputs on a given axis.

## 6. Spectral ratio heatmap

In the heatmap panel we divide MRS data elements into tiers of visual priority (**R3–R5**):

**Tier 1** Quantified spectral data

**Tier 2** Derived spectral data

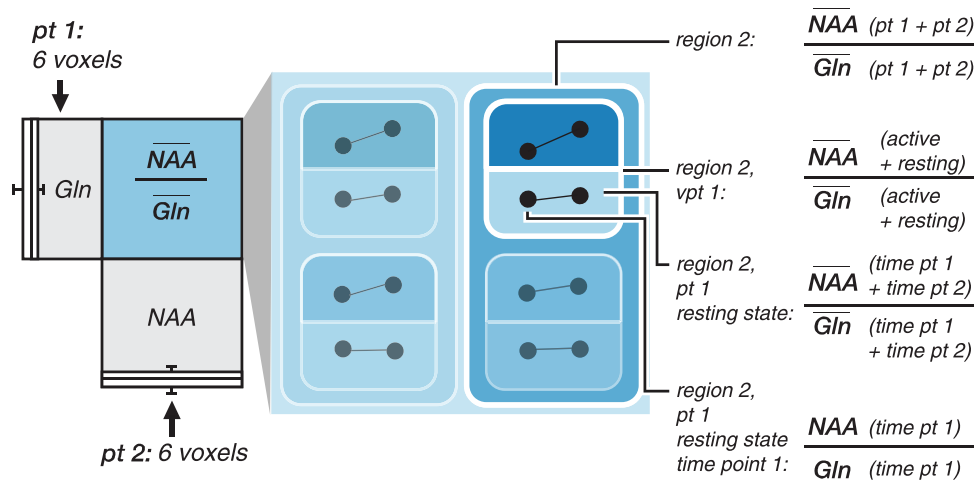
**Tier 3** Spectral metadata

**Tier 1** has primary importance; it consists of relative metabolite concentrations which are the result of pre-processing and quantification steps from the raw spectral acquisition. **Tier 2** comprises the complete set of metabolite ratios. It is used for comparison between user-defined groups as well as the following key attributes: spatial region, individual, brain activity state, and time point. Spatial region indicates the voxel sample position within the brain. Individual refers to a given patient included in the analysis. Time indicates either the number of separate spectral acquisitions performed on an individual over a study period, as in a longitudinal study, or recorded metabolite values within an acquisition session, as in a time-resolved MRS study. Finally, brain activity state indicates if the subject was in an *active* (task-explicit) state or *resting* (task-negative) state during signal acquisition. **Tier 3** includes metadata important for context and selection that are unnecessary to include as explicit encodings in the visualization: gender, age, and acquisition settings can have varying impact on the resulting concentrations and ratios of metabolites [25,36].

*Tier 1 encoding:* Visual perception research has shown that encoding position along a common axis is the most effective visual

channel for communicating quantitative information [37]. Box plots are a simple, ubiquitous and descriptive means of visually encoding statistical information about a dataset [38]. Since each MRS spectrum is essentially a multivariate set, where each metabolite is a variable, each metabolite in the spectrum then is tied to its own set of unique statistical information. We chose box plots over violin [39] or summary plots [40] to visualize **tier 1** data, as our goal with this tier is to provide clean, quickly readable insight to the input value range. Our use of box plots is additionally inspired by Blumenschein et al. [41], who used bars to encode aggregate dimensions in their work on table visualization. Bars and box plots are additionally well-recognized and easy to interpret; use of elements that were familiar to our target user group was an important design consideration. Furthermore, since box plots are only applicable when a dataset consists of five or more members, we introduce three variations depending on the number of inputs as illustrated in Fig. 4. For any of these variations, we first flatten the voxel hierarchy described in the spatial overview panel, and split the data into one voxel array per axis. In each array, we calculate the mean for each metabolite. In the case of a single spectral input, we use design variation A, which utilizes bars only, where height encodes the concentration of each metabolite (Fig. 4(A)). We calculate median, minimum, and maximum for two or more metabolite values on an axis. This corresponds to variation B, where height encodes the median value and whiskers encode the minimum and maximum metabolite concentration value, respectively (Fig. 4(B)). For five or more metabolites on an axis we additionally calculate the interquartile range. The box and whisker plot in variation C is utilized in this case, and shows the median, first and third quartiles, and the minimum and maximum value (Fig. 4(C)).

*Tier 2 encoding: Overview.* In **tier 2**, we visualize ratios between the mean along the *x*- and *y*-axes in a heatmap matrix (**R5**), as shown in Fig. 3(E). This effectively trades the low spatial resolution of MRS data for abstract resolution, focusing on biochemical concentrations in detail for a small region of interest. Each cell shows the aggregate ratio of the metabolite on the *x*-axis position to the corresponding *y*-axis metabolite, for instance mean Glutamine (Gln)/mean N-acetylaspartate (NAA), as illustrated in Fig. 5. We map the ratio value to a diverging red-blue colormap [42] inside each heatmap cell, as this color scheme is a familiar sight to our collaborators. In instances where the ratio is less than 1, we invert the ratio and switch the sign. To obtain a cleanly symmetric, divergent mapping structure we drop all values by 1 so the



**Fig. 5.** In nested ratio calculations, the cell background (A) is first mapped to color based the average of input metabolites on the *x*-axis divided by the average of input metabolites on the *y*-axis. Within the cell (B), the value of each input metabolite for all patients, at all time and brain state collections, is averaged and compared as a ratio for each spatial region. Within a spatial region, the average of each metabolite is compared for each patient, then for the brain state of each patient. The innermost step, time, takes a single metabolite input for both the numerator and denominator.

Case	region	individ.	state	time pt	visual	Case	region	individ.	state	time pt	visual
1	single	single	single	single		9	single	multiple	single	single	
2	single	single	dual	single		10	single	multiple	dual	single	
3	single	single	single	multiple		11	single	multiple	single	multiple	
4	single	single	dual	multiple		12	single	multiple	dual	multiple	
5	multiple	single	single	single		13	multiple	multiple	single	single	
6	multiple	single	dual	single		14	multiple	multiple	dual	single	
7	multiple	single	single	multiple		15	multiple	multiple	single	multiple	
8	multiple	single	dual	multiple		16	multiple	multiple	dual	multiple	

**Fig. 6.** Key tier 2 visual attributes include: brain spatial region, individual, brain state, and time point. We assign a unique glyph to each of these four attributes. Brain state is defined as active or resting; in absence of a classification we assume a resting state. All remaining three attributes may have single or multiple recordings. This produces 16 possible scenarios for spectral analysis. A sample visual is included for each scenario.

diagonal of the heatmap matrix is 0, rather than 1. Our aim is to draw attention to large input differences; this was identified as important for spectroscopy researchers. Red indicates a higher  $x$ -axis metabolite input while blue indicates a higher  $y$ -axis metabolite input. Equivalent inputs map to white. If an input value is 0, we map the cell color to dark grey. We originally thought to exclude such values from the heatmap, but on further discussion with our collaborators felt these were useful to include in order to preserve context. This heatmap view provides a means to visualize otherwise undetectable patterns in a rapid overview. To aid color interpretation and perception, our application includes a colormap legend to the right of the heatmap (Fig. 3(F)).

**Tier 2 encoding: Attributes.** Through a series of group interviews and individual shadowing sessions to the MR scanners we identified that, following an overview of all aggregated metabolite ratios, researchers are most interested in *comparing and summarizing* (T4) individual metabolites. For a given metabolite ratio, researchers first are interested in comparing *brain spatial regions*, as this can provide the most context for understanding ratio differences, e.g., in a tumor cohort study where voxels are acquired in the tumor region and in a healthy region of the brain. With spatial context, researchers can easily compare ratios between *individuals*. Assessing *brain activity state* is then most relevant in the context of the individual. After comparing the difference in active vs. resting brain state for an individual, the researcher may review the difference in these values over a cohort. Similarly, *time points* are best assessed first within a given brain state, then between states of an individual, before comparing between individuals.

In order to support experts in better identifying unexpected source ratios in a study, they thus need to evaluate four key attributes: (1) *brain spatial region*, (2) *individual*, i.e., *patient*, (3) *brain activity state*, and (4) *time point*. Furthermore, through each of these analysis stages we found that researchers prefer to maintain context between attributes to better understand sources of variation. This helped drive our development of a detailed metabolite ratio view that nests within each heatmap cell. Many MRS studies, particularly proof-of-concept research studies, by our collaborators often include around 20 subjects. They may sample up to four brain regions (although two is more typical), include up to three time points, i.e., pre-operative, post-operative, and long-term follow-up, and measure either a single or dual brain activity state. This space of attributes and approximate study size produces a set of 16 possible case scenarios to account for in our detailed comparison view.

**Tier 2 encoding: Detail.** Given the low number of key attributes, we found a simple glyph representing each attribute to be the most conducive to user analysis. Our glyph choice and design was mainly inspired by findings from unit visualization research, mainly the ATOM grammar by Park et al. [20], for this method's demonstrated strong intuition and interaction properties. Since our target study sizes are typically relatively small, we avoid issues with display and perceptual scalability from which unit visualizations often suffer. To maintain important context in the analysis flow, we nest glyphs to mirror the order of analysis preferred by researchers. Our glyph nesting design was inspired by dimensional stacking visualization techniques pioneered in XmdvTool

and N-land by Ward et al. [43,44]. Since nested glyphs can form complex shapes, we chose glyphs that were simple and familiar to our collaborators to reduce interpretation difficulties. Although we discussed different stroke styles for glyphs, for simplicity and clarity our ultimate design uses a solid hairline stroke for each of the four attributes. Experts felt that changes in stroke weight or style was distracting and overemphasized elements; this may bias conclusions.

The visual design for this detailed view is mapped from a series of nested ratios. Inside each cell we flatten the data to a single voxel array, skipping any duplicate voxels. We then determine ratios for each of the key attributes, where available, in a nested fashion that mirrors the preferred order of user analysis: the ratio for each spatial region (using the average of all individuals for this region), each individual (using the average of all states for the given individual in a given region), each state (using the average of all time points for a given state of an individual from a given region) and each time point, as shown in Fig. 5(B). These nested values then map to the appropriate glyph.

We represent spatial regions as rounded rectangular glyphs. We chose rounded corners to distinguish spatial glyphs from the square shape of the heatmap overview cell. Furthermore, the rounded corners leave space to reveal the heatmap cell color, thereby subtly preserving context within the detail view. In each cell, we evenly divide the space vertically by the number of distinct regions sampled. Individuals are presented as filled disks when only shown in a single time acquisition (e.g., case 9), expanding to rounded squares when time series data are incorporated (e.g., case 3). This shape change permits a spark line to move evenly across the space without going outside the border of the enclosing glyph. Shapes scale to fill space within their frame. In instances where different brain activity states are analyzed, we divide the shape in half horizontally (e.g., case 2). This feature was important to include for our collaborators who perform time-resolved spectroscopy, as this is not available in other tools. Finally, we encode different time acquisitions as points connected via a spark line, inspired by Meyer et al. in their work, Pathline [45]. This spark line is nested into the relevant glyph: if a multi time step series is captured in a study analyzing different brain states, the spark line is placed within each state half-moon glyph (e.g., case 4). If analysis is only for a single activity state, the spark line nests inside the individual glyph (e.g., cases 3, 11), or inside spatial glyphs for a single patient (e.g., case 7). The remaining cases comprise different permutations of these spatial region, individual, brain state, and time point arrangements.

For example, consider an instance of scenario 16: two patients are sampled in two regions of the brain four times in a year. During two acquisition times the subjects were asked to perform a task (active brain state), while the other two times were asked to relax (resting brain state). This produces a total of 12 unique measurements, 6 per patient. The overview cell is calculated by averaging the 6 values of Gln for patient 1 and the 6 values of NAA for patient 2, and dividing the result of NAA into Gln. Inside the cell, we compute this ratio as a series of nested averages for each of the four key attributes, as depicted in Fig. 5: (1) spatial region, (2) individual, (3) brain state, and (4) time point. For each, we average the metabolite concentrations before computing the ratio. For additional detail view images and example tasks of each scenario in a more complex dataset, we refer interested readers to the supplementary material *SpectraMosaic Detail Case Scenarios*.

Hovering facilities display the ratio value for each cell or attribute of interest (R5). Displaying this numerical value provides a safeguard against possible distortions of color perceptions that may occur with our chosen glyph nesting structure. This value is displayed in red text if one metabolite input exhibits an uncertainty above 15% (Cramér-Rao lower bound) [46]. This information may

be used to assess both the quality of the measurement and the accuracy of the spectral processing and quantification steps.

**Tier 3** of MRS data consists of metadata information used for context and selection. We depict this information in a table below the heatmap. Gender, age, and echo time comprise other important patient attributes to track because the shape of the spectrum can vary considerably with these factors—for example, the lactate peak is virtually undetectable in healthy babies [36], but is nearly always measurable in healthy adults with increased neural activation [25]. Acquisition settings are also important, as different echo times will yield a vastly different spectral representation for the same patient.

## 7. Implementation

SpectraMosaic is a web-based application implemented with HTML, CSS, Javascript, as well as the D3 [47], P5, and gridster Javascript libraries. It was developed as a web application to allow for easy integration and use within the hospital network (R1). A Python back end integrates MATLAB [33] and Tarquin [27] components in the preprocessing steps. Assets are stored on the client and fetched on-demand. Our visualization tool code is open source and is publicly available at <https://github.com/mmiv-center/spectramosaic-public>.

## 8. Case study


We evaluated the utility of SpectraMosaic as a research tool using a giardiasis MRS case study. Giardiasis is a parasite-borne disease affecting the small intestine caused by drinking water contamination. The metabolic byproducts of this disease are subtle, but have been shown to be detectable by MRS [48,49]. The goal of this study is to explore and identify possible metabolic indicators for infection using our tool.

Collected in Bergen, Norway, study data comprised two patients imaged some months apart in three different regions of the brain at a single echo time (TE 35 ms). For one region (prefrontal region) two different TE parameter settings were used (TE 35 ms and TE 144 ms). These data were analyzed by three volunteers recruited from the fMRI/MRS research group in Bergen. All three provided feedback on earlier interfaces of the SpectraMosaic application, and are not co-authors of this work. User A is an MR physicist specializing in development and refinement of spectroscopy protocols for clinical studies of neuropsychiatric and developmental disorders. User B, also an MR physicist, uses <sup>31</sup>P- and <sup>13</sup>C-labeled pyruvate timecourse data to study real-time metabolism. User C is a cognitive neuroscientist who uses MRS in conjunction with fMRI in research on neurodegenerative and developmental disorders, e.g., Parkinson's disease, stroke, and stuttering. We processed the data in advance to focus evaluation on the visual web tool; this step included de-identification of patient-specific information.

*Case workflow feedback.* After a brief introduction to the tool, users analyzed this case following a “think-aloud” protocol [50]. We conducted follow-up interviews after the analysis was complete, which we summarize and discuss in this section.

All three users began with an overview of the spectral graphs and voxel position for each imaged brain region (T1). User A investigated spectral graphs by region, while B and C explored by patient. Users A and B commented that this overview provides an important quality assurance check for each acquisition. Since all three users are familiar with MRS, they agreed with our decision to exclude labeling of spectral peaks; they felt this would have been unnecessary and distracting to include. All commented that the hippocampus region spectra looked strange, which could be indicative of either a pathology or acquisition problem. They noted that this region is particularly difficult to image well, and requires deeper investigation.





**Fig. 7.** Heatmap inspection in a two-patient, multi-voxel study acquired at two TEs: 35 ms (x-axis) and 144 ms (y-axis). Investigating the Alanine (Ala)/Alanine (Ala) cell reveals a higher measurement in the TE 144 ms group. However, the tooltip indicates that this ratio may be unreliable due to a poor model fit.

All users then explored available group presets and experimented with creation of custom groups for analysis (T2). They agreed that the presets particularly improved the practical usability of the tool, stating that these were comprehensive and largely removed the need to make custom groups. All users experimented with adding a subset of basis set metabolites (Fig. 3(D)) to the heatmap view, although they felt that analysis of all metabolites is a useful first step for exploring new hypotheses. However, they agreed that subset metabolite analysis is useful as hypotheses are refined to a narrower metabolite set.

Feedback was positive for the alphabetical ordering of metabolites on heatmap axes. User C strongly felt that any statistics-based ordering method would make interpretation too difficult because they would spend too much time locating metabolites along the axes. All users agreed that the representation of metabolite relative concentrations as whisker bar or box plots was extremely useful, as it offered additional insight into unexpected values observed in the heatmap. User B stated: “Checking the range on the metabolite inputs helps me as a first check; a huge range could indicate a [brain region] area effect or a bad acquisition. I can easily then verify this by checking the spectral graph in the other panel.” All users noted a massive range for Gamma-Aminobutyric acid (GABA) in this study, and were able to quickly conclude that the acquisition technique used is not effective for this metabolite. For this study and others acquired on the same scanner, through the same technique and parameters, this representation allows for a straightforward relative comparison of metabolites before ratio computation (T3).

In the spectral ratio heatmap, user A was primarily interested in exploring ratios at different echo times (TE) (T4). We see this exploration in Fig. 7; TE 144 ms voxels are placed on the y-axis while TE 35 ms voxels are placed on the x-axis. This user focused on the diagonal of the matrix, and primarily on examining known metabolites implicated in giardia infection, e.g., Alanine. Although this ratio shows relative similarity, we note that the model fit for this metabolite is outside the accepted range. This requires further investigation. User B also compared different echo times, but over the entire matrix space for any unexpected dark color regions. For each unexpected cell, the user noted whether this could be pathology, or an acquisition problem.

All three users were also interested in comparing ratios of metabolites between patients for each of the three measured brain regions (T4). They first filtered out TE 144 ms acquisitions, then

arrayed each patient on opposing axes. Assuming both patients are healthy, we would expect that the patient glyphs for all spatial regions would show similar values. All three users noted an unexpected, relatively large difference for Lactate/Total Creatine in the hippocampal region (Fig. 8). To investigate this disparity, users A and B first verified whether the value met the threshold for each patient. The value did not meet the threshold for the female patient, indicating an unreliable fit. Users then reviewed the spectral graph of the hippocampal acquisition for this patient, noting its abnormality—users concluded that this merits deeper investigation, and likely requires a new acquisition.

*Summary feedback.* All three participants felt that SpectraMosaic was useful and could augment their standard workflow for deeper insights into spectral data. User A noted that the visual feedback on the model fit in each ratio provides invaluable data quality information. User B stated, “The linking between the glyph ratios, the spectra, the table, and the images is incredibly useful for us—whenever we look at metabolite results we always want to go back to the raw spectra and see if this makes sense, and if the quality is good, and this makes it really easy to see. I see this tool as being useful to verify assumptions I have going in to the study, and to explore the entire range for quality checks that might affect the results that I’m expecting.”

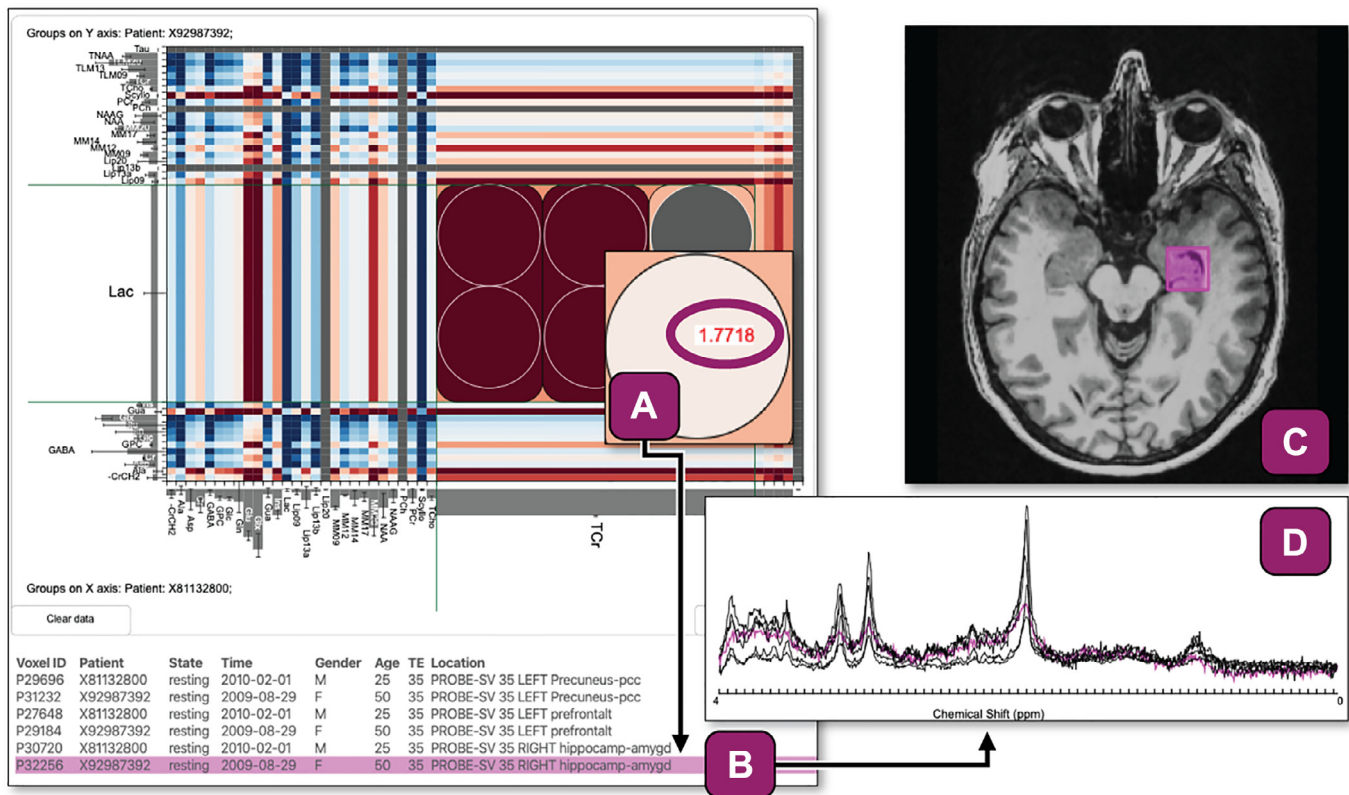
All participants felt the nested glyphs were integral elements of the metabolite ratio exploration process. The detail glyph view provided a means to quickly drill into an unexpected ratio and identify the possible source(s), while easily retaining contextual information from the surrounding heatmap. User C noted: “This [spectral heatmap] overview and detail glyph feature is useful to have a closer look at, for example, neurodegeneration [in Parkinson’s] with the loss of dopaminergic connections, as seen with concentrations of glutamate or GABA... and it is ideal for testing new protocols against established protocols.” Furthermore, experts agreed that the glyph design and nesting structure was intuitive and clear in all case scenarios, even in larger, more complex studies. All three stated that interpreting these glyphs was not difficult, particularly when compared to the very steep learning curve to interpreting spectral data through their standard approach. They agreed that the inclusion of the legend was helpful when first familiarizing themselves with the system, but that they had little need to reference it after the first few minutes of heatmap exploration. However, two experts commented that our mapping of vertical time points could be scaled differently to more clearly demonstrate relative ratio value changes, which were at times difficult to recognize. For detailed expert feedback on clarity and interpretability of the nested glyph structure in all 16 possible case scenarios in a larger dataset, we refer interested readers to *SpectraMosaic Detail Case Scenarios* in the supplementary material.

All users indicated interest in an option to extract spectral heatmap visuals and data for subsequent statistical analysis; user A expressed interest in seeing this output to the hospital PACS for access by radiologists to aid in more rapid interpretation of spectroscopy data for more widespread clinical use.

## 9. Discussion and limitations

In the case evaluation of SpectraMosaic we found that our tool provides new, interesting insights on metabolic profiles at different aggregation levels.

Our task analysis showed that experts were particularly interested in large metabolite differences. Although our diverging color mapping approach in the heatmap is effective in demonstrating large differences between metabolites, subtle differences are less obvious. Investigation into fine grained color mapping options or user-defined color map scaling may help more clearly highlight these instances. This extends to our plotting of time points, where



**Fig. 8.** Following overview of the metabolite ratios between the individual patients, the user inspects Lactate (Lac)/Total creatine (TCr) ratio between two patients for all three regions at TE 35 ms. The user notes a high lactate measurement for the female patient relative to the other measured regions (A). Subsequent inspection of spectral metadata (B), the spectral graph (C), and the brain region in which this measurement was acquired (D) help establish reasons for this difference.

subtle ratio changes could benefit from a logarithmic axis scaling approach to highlight such changes to users.

While our decision to sort metabolite inputs in a consistent order limits the ability for pattern recognition *within* a study, this approach allows for pattern recognition *between* studies, where users can begin to observe a typical “footprint” for certain acquisition techniques.

Although this is uncommon for our collaborators, we also note that if data are not acquired from the same scanner and same parameters, the utility of the bar and box encoding becomes more limited. This is because different scanners and different parameters can vastly change the metabolite concentrations; in this case the ratio heatmap becomes the primary tool for comparative analysis.

Our visual design, particularly with reference to the nested glyphs in the detail view, was guided by collaborative discussions with research experts. These relatively small study sizes are conducive to nested unit visualizations, and in this iteration of the application were not designed to scale to, e.g., hundreds of patients. With respect to the scalability of groupings within our planned design, we conducted a preliminary assessment of nested glyph interpretability for each case scenario using a larger study. We provide the results of this assessment in the supplementary material (*SpectraMosaic Detail Case Scenario*). Our collaborators even indicated that they could envision this approach scaling beyond 20 patients for some scenarios. Additionally, we could incorporate clustering or an additional design layer for automatically- or user-generated groupings for further scalability (T4).

Lastly, while our glyph system covers all main use cases, we found that echo time is varied in research studies more often than initially expected. This frequency of use may imply that this attribute should be encoded at the second priority visualization tier,

rather than its current third level. However, comparison of different echo times beyond an overview level is of less clinical interest than the four attributes we have discussed. Inclusion of a fifth glyph would require careful consideration.

## 10. Conclusions and future work

In this design study we contributed a characterization of the data, task, and design requirements for the development of SpectraMosaic, followed with an expanded tiered visual encoding system and pipeline. We performed case studies with three domain experts to validate our tool in spectroscopy clinical research and protocol development. MR spectroscopy is a ripe area for continued visualization research.

The flexible design of our tool allows for a number of possible extensions; this may include investigation into additional statistical measures relevant for comparative analysis, e.g., correlation. Although this paper focuses on  $^1\text{H}$ -MRS,  $^{31}\text{P}$ -MRS and  $^{23}\text{Na}$ -MRS analysis may also be integrated to our tool. While we offer basic mechanisms for uncertainty visualization, exploring additional means for uncertainty feedback in the heatmap cells and glyphs can offer deeper insights into the data. Finally, although typical MRS cohort studies are relatively small, exploration of methods to extend our visual encoding system to successfully manage larger cohorts may further increase tool usability.

Automatic interface adjustment based on acquisition technique offers a valuable investigation of parameter space analysis in MRS. Exploration of the most salient features to reveal for, e.g., PRESS versus MEGA-PRESS, may help experts more effectively identify interesting ratios for further investigation. Beyond the medical domain, an additional interesting line of inquiry would be to explore the adaptability of our abstracted line tasks paired with our visual en-

coding system in other areas facing similar challenges with heterogeneous multidimensional data, such as meteorology or geophysics.

### Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### CRediT authorship contribution statement

**Laura Garrison:** Conceptualization, Methodology, Investigation, Visualization, Writing - original draft. **Jakub Vašiček:** Methodology, Software. **Alexander R. Craven:** Resources, Software, Validation, Data curation, Writing - review & editing. **Renate Grüner:** Resources, Supervision, Writing - review & editing. **Noeska N. Smit:** Writing - review & editing. **Stefan Bruckner:** Conceptualization, Methodology, Supervision, Writing - review & editing, Funding acquisition.

### Acknowledgments

We thank the UiB VisGroup and our collaborators at the Mohn Medical Imaging and Visualization Centre (MMIV) and Haukeland University Hospital: Lars Erslund, Karsten Specht, Frank Riemer, Hauke Bartsch, Kenneth Hugdahl, and Gerard Dwyer for their invaluable insights and feedback. This work is supported by the University of Bergen and the Trond Mohn Foundation in Bergen, Norway (#813558, Visualizing Data Science for Large Scale Hypothesis Management in Imaging Biomarker Discovery; and #811255).

### Supplementary material

As supplementary material, we provide a video demonstrating the SpectraMosaic application features and workflow. We additionally include *SpectraMosaic Detail Case Scenarios*, which detail each of the 16 detail case scenarios in a more complex case study.

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.cag.2020.08.001.

### References

- [1] Van Der Graaf M. In vivo magnetic resonance spectroscopy: basic methodology and clinical applications. *Eur Biophys J* 2010;39(4):527–40. doi:10.1002/9780470882221.
- [2] Swanberg KM, Landheer K, Pitt D, Juchem C. Quantifying the metabolic signature of multiple sclerosis by in vivo proton magnetic resonance spectroscopy: Current challenges and future outlook in the translation from proton signal to diagnostic biomarker. *Front Neurol* 2019;10. doi:10.3389/fneur.2019.01173.
- [3] Verma A, Kumar I, Verma N, Aggarwal P, Ojha R. Magnetic resonance spectroscopy – revisiting the biochemical and molecular milieu of brain tumors. *BBA Clin* 2016;5:170–8. doi:10.1016/j.bbacli.2016.04.002.
- [4] Provencher SW. Automatic quantitation of localized in vivo <sup>1</sup>H spectra with LCModel. *NMR Biomed* 2001;14(4):260–4. doi:10.1002/nbm.698.
- [5] Garrison L, Vašiček J, Grüner R, Smit N, Bruckner S. Spectramosaic: an exploratory tool for the interactive visual analysis of magnetic resonance spectroscopy data. In: Proceedings of ninth EG workshop on visual computing for biology and medicine (VCBM); 2019.
- [6] Wilson M, Reynolds G, Kauppinen RA, Arvanitis TN, Peet AC. A constrained least-squares approach to the automated quantitation of in vivo <sup>1</sup>H magnetic resonance spectroscopy data. *Magn Reson Med* 2011a;65(1):1–12. doi:10.1002/mrm.22579.
- [7] Amir Khanov A, Fröhler B, Kastner J, Gröller E, Heinzl C. InSpectr: multi-modal exploration, visualization, and analysis of spectral data. *Comput Graph Forum* 2014;33(3):91–100. doi:10.1111/cgf.12365.
- [8] Gleicher M, Albers D, Walker R, Jusufi I, Hansen CD, Roberts JC. Visual comparison for information visualization. *Inf Visual* 2011;10(4):289–309. doi:10.1177/1473871611416549.
- [9] Bruckner S, Möller T. Isosurface similarity maps. *Comput Graph Forum* 2010;29(3):773–82. doi:10.1111/j.1467-8659.2009.01689.x.
- [10] Fröhler B, Amir Khanov A, Kastner J, Gröller E, Heinzl C. Multimodal visualization and analysis of spectral and XCT data. In: Proceedings of ninth Forschungsforum der sterreichischen Fachhochschulen; 2015.
- [11] Feng D, Lee Y, Kwock L, Taylor RM. Evaluation of glyph-based multivariate scalar volume visualization techniques. In: Proceedings of the Symposium on Applied Perception in Graphics and Visualization; 2009. p. 61–8. doi:10.1145/1620993.1621006.
- [12] Feng D, Kwock L, Lee Y, Taylor RM. Linked exploratory visualizations for uncertain MR spectroscopy data. *Visual Data Anal* 2010;7530:753004. doi:10.1117/12.839818.
- [13] Nunes M, Rowland B, Schlachter M, Ken S, Matkovic K, Laprie A, et al. 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. p. 53–62. doi:10.1109/VAST.2014.7042481.
- [14] Matkovic K, Freiler W, Gracanin D, Hauser H. ComVis: a coordinated multiple views system for prototyping new visualization technology. In: Proceedings of International Conference Information visualisation; 2008. p. 215–20. doi:10.1109/IV.2008.87.
- [15] Wolf I, Vetter M, Wegner I, Böttger T, Nolden M, Schöbinger M, et al. The medical imaging interaction toolkit. *Med Image Anal* 2005;9(6):594–604. doi:10.1016/j.media.2005.04.005.
- [16] Marino J, Kaufman A. Prostate cancer visualization from MR imagery and MR spectroscopy. *Comput Graph Forum* 2011;30(3):1051–60. doi:10.1111/j.1467-8659.2011.01954.x.
- [17] Jawad M, Molchanov V, Linsen L. Coordinated image- and feature-space visualization for interactive magnetic resonance spectroscopy imaging data analysis. In: Proceedings of the fourteenth international joint conference on computer vision, imaging and computer graphics theory and applications – volume 3 IVAPP: IVAPP; 2019a. p. 118–28. doi:10.5220/0007571801180128.
- [18] Jawad M, Evers M, Gerwing A, Herick M, Seibert D, Bauer J, et al. A visual analytics approach for comparing cohorts in single-voxel magnetic resonance spectroscopy data. *Biomed Visual* 2019b;115–36. doi:10.1007/978-3-030-14227-8\_9.
- [19] Bertin J, Berg WJ, Wainer H. *Semiotics of graphics: diagrams, networks, maps*. University of Wisconsin press Madison; 1983. doi:10.1080/00690805.1987.10438353.
- [20] Park D, Drucker S, Fernandez R, Elmquist N. Atom: A grammar for unit visualizations. *IEEE Trans Visual Comput Graph* 2017;24(12):3032–43. doi:10.1109/TVCG.2017.2785807.
- [21] Stolte C, Tang D, Hanrahan P. Polaris: a system for query, analysis, and visualization of multidimensional relational databases. *IEEE Trans Visual Comput Graph* 2002;8(1):52–65. doi:10.1145/1400214.1400234.
- [22] Klemm P, Oeltze-Jafra S, Lawonn K, Hegenscheid K, Völzke H, Preim B. Interactive visual analysis of image-centric cohort study data. *IEEE Trans Visual Comput Graph* 2014;20(12):1673–82. doi:10.1109/tvcg.2014.2346591.
- [23] Hellström J, Zapata RR, Libard S, Wikström J, Ortiz-Nieto F, Alafuzoff I, et al. The value of magnetic resonance spectroscopy as a supplement to MRI of the brain in a clinical setting. *PLoS One* 2018;13(11). doi:10.1371/journal.pone.0207336.
- [24] Ulmer S, Backens M, Ahlhelm FJ. Basic principles and clinical applications of magnetic resonance spectroscopy in neuroimaging. *J Comput Assist Tomogr* 2016;40(1):1–13. doi:10.1097/RCT.0000000000000322.
- [25] Stagg C, Rothman DL. *Magnetic resonance spectroscopy: tools for neuroscience research and emerging clinical applications*. Academic Press; 2013. doi:10.1016/C2011-0-09647-3.
- [26] Stefan D, Di Cesare F, Andrasescu A, Popa E, Lazariev A, Vescovo E, et al. Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package. *Measur Sci Technol* 2009;20(10):104035. doi:10.1088/0957-0233/20/10/104035.
- [27] Wilson M, Reynolds G, Kauppinen RA, Arvanitis TN, Peet A. A constrained least-squares approach to the automated quantitation of in vivo <sup>1</sup>H magnetic resonance spectroscopy data. *Magn Reson Med* 2011b;65(1):1–12. doi:10.1002/mrm.22579.
- [28] Crane JC, Olson MP, Nelson SJ. SIVIC: open-source, standards-based software for DICOM MR spectroscopy workflows. *J Biomed Imaging* 2013;2013:12. doi:10.1155/2013/169526.
- [29] Purvis LAB, Clarke WT, Biasioli L, Valkovič L, Robson MD, Rodgers CT. OXSA: An open-source magnetic resonance spectroscopy analysis toolbox in MATLAB. *PLoS One* 2017;12(9):e0185356. doi:10.1371/journal.pone.0185356.
- [30] Mullins PG, McGonigle DJ, O’Gorman RL, Puts NA, Vidyasagar R, Evans CJ, et al. Current practice in the use of MEGA-PRESS spectroscopy for the detection of GABA. *Neuroimage* 2014;86(1):43–52. doi:10.1016/j.neuroimage.2012.12.004.
- [31] Brehmer M, Munzner T. A multi-level typology of abstract visualization tasks. *IEEE Trans Visual Comput Graph* 2013;19(12):2376–85. doi:10.1109/TVCG.2013.124.
- [32] Madhu B, Jauhainen A, McGuire S, Griffiths JR. Exploration of human brain tumour metabolism using pairwise metabolite-metabolite correlation analysis (MMCA) of HR-MAS <sup>1</sup>H NMR spectra. *PLoS One* 2017;12(10):e0185980.
- [33] MATLAB. 9.7.0.1190202 (R2019b). Natick, Massachusetts: The MathWorks Inc.; 2018.
- [34] Corey RB, Pauling L. Molecular models of amino acids, peptides, and proteins. *Rev Sci Instrum* 1953;24(8):621–7. doi:10.1063/1.1770803.
- [35] Perin C, Dragicevic P, Fekete J. Bertifier: New interactions for crafting tabular visualizations. *IEEE Trans Visual Comput Graph* 2014;20(12):2082–91. doi:10.1109/TVCG.2014.2346279.
- [36] Xu D, Vigneron D. Magnetic resonance spectroscopy imaging of the newborn brain—a technical review. *Semin Perinatol* 2010;34(1):20–7. doi:10.1053/j.semperi.2009.10.003.

- [37] Cleveland WS, McGill R. Graphical perception: theory, experimentation, and application to the development of graphical methods. *J Am stat Assoc* 1984;79(387):531–54. doi:[10.2307/2288400](https://doi.org/10.2307/2288400).
- [38] Williamson DF, Parker RA, Kendrick JS. The box plot: a simple visual method to interpret data. *Ann Intern Med* 1989;110(11):916–21. doi:[10.7326/0003-4819-110-11-916](https://doi.org/10.7326/0003-4819-110-11-916).
- [39] Hintze JL, Nelson RD. Violin plots: a box plot-density trace synergism. *Am Stat* 1998;52(2):181–4. doi:[10.1080/00031305.1998.10480559](https://doi.org/10.1080/00031305.1998.10480559).
- [40] Potter K, Kniss J, Riesenfeld R, Johnson CR. Visualizing summary statistics and uncertainty, 29; 2010. p. 823–32. doi:[10.1111/j.1467-8659.2009.01677.x](https://doi.org/10.1111/j.1467-8659.2009.01677.x).
- [41] Blumenschein M, Behrisch M, Schmid S, Butscher S, Wahl DR, Villinger K, et al. Smartexplore: simplifying high-dimensional data analysis through a table-based visual analytics approach. In: *Proceeding of 2018 IEEE conference on visual analytics science and technology (VAST)*; 2018. p. 36–47. doi:[10.1109/VAST.2018.8802486](https://doi.org/10.1109/VAST.2018.8802486).
- [42] Harrower M, Brewer CA. ColorBrewer.org: an online tool for selecting colour schemes for maps. *Cartogr J* 2003;40(1):27–37. doi:[10.4324/9781351191234-18](https://doi.org/10.4324/9781351191234-18).
- [43] Ward MO. Xmdvtool: integrating multiple methods for visualizing multivariate data. In: *Proceedings of the conference on visualization*. IEEE Computer Society Press; 1994. p. 326–33. doi:[10.1109/visual.1994.346302](https://doi.org/10.1109/visual.1994.346302).
- [44] Ward M, LeBlanc JT, Tipnis R. N-land: a graphical tool for exploring n-dimensional data. In: *Proceedings of computer graphics international conference*; 1994. p. 131–41.
- [45] Meyer M, Wong B, Styczynski M, Munzner T, Pfister H. Pathline: A tool for comparative functional genomics. *Comput Graph Forum* 2010;29(3):1043–52. doi:[10.1111/j.1467-8659.2009.01710.x](https://doi.org/10.1111/j.1467-8659.2009.01710.x).
- [46] Cavassila S, Deval S, Huegen C, Van Ormondt D, Graveron-Demilly D. Cramér-Rao bounds: an evaluation tool for quantitation. *NMR Biomed: Int Journal Devot Dev Appl Magn Resonanc In Vivo* 2001;14(4):278–83. doi:[10.1002/nbm.701](https://doi.org/10.1002/nbm.701).
- [47] Bostock M, Ogievetsky V, Heer J. D3 data-driven documents. *IEEE Trans Visual Comput Graph* 2011;17(12):2301–9. doi:[10.1109/TVCG.2011.185](https://doi.org/10.1109/TVCG.2011.185).
- [48] Edwards MR, Gilroy FV, Jimenez BM, O'Sullivan WJ. Alanine is a major end product of metabolism by giardia lamblia: a proton nuclear magnetic resonance study. *Mol Biochem. Parasitol.* 1989;37(1):19–26. doi:[10.1016/0166-6851\(89\)90098-4](https://doi.org/10.1016/0166-6851(89)90098-4).
- [49] Vermathen M, Müller J, Furrer J, Müller N, Vermathen P. 1H HR-MAS NMR spectroscopy to study the metabolome of the protozoan parasite giardia lamblia. *Talanta* 2018;188:429–41. doi:[10.1016/j.talanta.2018.06.006](https://doi.org/10.1016/j.talanta.2018.06.006).
- [50] Lewis C. Using the "thinking-aloud" method in cognitive interface design. IBM TJ Watson Research Center Yorktown Heights, NY; 1982. doi:[10.1080/00140139008927157](https://doi.org/10.1080/00140139008927157).