Coordinated Image- and Feature-space Visualization for Interactive Magnetic Resonance Spectroscopy Imaging Data Analysis

Muhammad Jawad, Vladimir Molchanov and Lars Linsen
Westfälische Wilhelms-Universität Münster, Germany

Keywords: Multidimensional Data Visualization, Medical Visualization, Coordinated Views, Spectral Imaging Analysis.

Abstract: Magnetic Resonance Spectroscopy Imaging (MRSI) is a medical imaging method that measures per voxel a spectrum of signal intensities. It allows for the analysis of chemical compositions within the scanned tissue, which is particularly useful for tumor classification and measuring its infiltration of healthy tissue. Common analysis approaches consider one metabolite concentration at a time to produce intensity maps in the image space, which does not consider all relevant information at hand. We propose a system that uses coordinated views between image-space visualizations and visual representations of the spectral (or feature) space. Coordinated interaction allows for analyzing both aspects and relating the analysis results back to the other for further investigations. We demonstrate how our system can be used to analyze brain tumors.

1 INTRODUCTION

Magnetic Resonance Spectroscopy Imaging (MRSI) is an in-vivo medical imaging method for measuring chemical compositions of scanned tissues. The compositions in the form of metabolite concentrations can be computed from spectral information of chemical resonance (Gujar et al., 2005). While T1- or T2-weighted Magnetic Resonance Imaging (MRI) data allows for the detection of tumors and to determine their shape, MRSI provides additional information on their metabolic activity. Such information, on the one hand, allows for a classification of the tumor with respect to its malignancy (from benign to malignant), and, on the other hand, for an investigation, whether the tumor has already started infiltrating surrounding "healthy" tissue (Burnet et al., 2004).

The metabolic information is extracted from the spectrum for each voxel of the MRSI data in a preprocessing step, see Section 3.2. Afterwards, the information at hand is the metabolic concentration of all extracted metabolites per voxel. The set of metabolites form a multidimensional space, which we refer to as the spectral or feature space. Each voxel is reflected by a point in this multidimensional space. On the other hand, the voxels have a spatial arrangement in the image space.

Common analysis tools for MRSI data in clinical settings visualize the spatial distribution of individual metabolite concentrations using color mapping of an image slice. Thus, only a single metabolite is investigated at a time. Therefore, a lot of information is being neglected and the interplay of metabolites concentrations cannot be analyzed. We propose a novel tool that integrates all information at hand for a comprehensive analysis of MRSI data. Our approach is based on coordinated views of visualization in image and feature space. For image space, we also use slice-based visualizations, as they avoid occlusion and it is common to only scan a few slices in MRSI. MRSI visualizations are overlaid with MR images and combined with automatic MRI segmentation results. For feature space, multidimensional data visualization methods are used. Given the image segmentation result, the multidimensional data are labeled accordingly and we apply interaction methods to separate the labeled classes. Using star-coordinates encoding, the separations can be related back to the individual dimension, basically allowing for conclusions, which metabolite concentration allow for class separation. The methodology is detailed in Section 4.

In Section 5, we apply our methods to a scenario for MRSI data analysis to investigate brain tumors. In 2012, WHO reported 256,213 brain cancer cases, of which 189,382 deaths have been recorded (Board, PDQ Adult Treatment Editorial, 2018). We show how our coordinated views allow for analyzing the brain tumor's chemical composition as well as the infiltration of surrounding tissue that when based on MRI data may be classified as non-tumor region.

Jawad, M., Molchanov, V. and Linsen, L.
Coordinated Image- and Feature-space Visualization for Interactive Magnetic Resonance Spectroscopy Imaging Data Analysis.
DOI: 10.5220/0007571801180128
In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2019)*, pages 118-128
ISBN: 978-989-758-354-4

2 RELATED WORK

Nunes et al. (Nunes et al., 2014a) provided a survey of existing methods for analyzing MRSI data. They conclude that no approach exists that analyzes all metabolic information. They report that MRSI data visualization packages provided with commercial scanners such as SyngoMR and SpectroView only provide color mapping of individual metabolic concentrations (or a ratio of metabolic concentrations) in image space. Other tools such as Java-based Magnetic Resonance User Interface (jMRUI) (Stefan et al., 2009) and Spectroscopic Imaging Visualization and Computing (SIVIC) (Crane et al., 2013) that are also widely used in clinical practice provide a similarly restricted functionality.

Raschke et al. (Raschke et al., 2014) proposed an approach to differentiate between tumor and non-tumor regions by showing the relation of concentrations of two selected metabolites in a scatterplot. By plotting the regression line, they identified abnormal regions by data points that are far from the line. Similarly, Rowland et al. (Rowland et al., 2013) use scatterplots to inspect three different pairs of metabolites tumor analysis. These procedures only target selected metabolites, which are chosen a priori.

An approach to exploit the metabolic information better was presented by Maudsley et al. (Maudsley et al., 2006). They developed the Metabolite Imaging and Data Analysis System (MIDAS) for MRSI pre-processing and visualization. Users can view histograms of metabolites, but relations between metabolites cannot be studied. Feng et al. (Feng et al., 2010) presented the Scaled Data Driven Sphere (SDDS) technique, where information of multiple metabolite concentrations is combined in a glyph-based visualization using mappings to color and size. Obviously, the amount of dimensions that can be mapped is limited. They overlay the images with the glyphs and link them to parallel coordinate plots. Users can make selections in the parallel coordinate plot and observe respective spatial regions in the image space. As a follow-up of their survey, Nunes et al. (Nunes et al., 2014b) presented a system that couples the existing systems ComVis (Matkovic et al., 2008) and MITK (Wolf et al., 2004). They use scatterplot, histogram, and parallel coordinate plot visualizations for analyzing the metabolic features space. The visualizations provide linked interaction to image-space representations such that brushing on metabolic data triggers highlighting in image space.

We build upon the idea of using coordinated interactions in feature and image space, but enhance the functionality significantly: We incorporate segmentation results, provide means to separate the labeled data in feature space, allow for a comparative visualization of classes, and provide a single tool that integrates all information and allows for fully coordinated interaction in both directions.

3 BACKGROUND

In this section, we provide background on the imaging method, describe the executed pre-processing steps, describe the data at hand after imaging and pre-processing, and the driving questions.

3.1 Imaging

MRSI is an in-vivo medical imaging method, where per voxel a whole spectrum of intensities is recorded. While MRI only measures water intensities per voxel, MRSI measures intensities at different radio frequencies. Most commonly ^{1}H MRSI is used, where the signal of hydrogen protons (^{1}H) in different chemicals is measured in the form of intensity peaks. As this measurement is performed for different frequencies, different chemicals can be detected within each voxel. The intensity peaks within the frequency spectrum are quantified as parts per million (ppm).

Measuring entire spectra in MRSI comes at the expense of much lower spatial resolution when compared to MRI. Using 1.5T or 3T scanners yields voxel sizes of about $10mm \times 10mm \times 8mm$ (Nunes et al., 2014a; McKnight et al., 2001). 7T scanners yielding higher resolutions are barely used in clinical practice due to high costs (Scheenen et al., 2008). Typically, MR images are taken (using T1 or T2 relaxation) in addition to MRS images within the same session. The MR images provide anatomical information at a higher resolution.

3.2 Pre-processing

The main goal of the pre-processing step is to quantify the various chemicals, referred to as metabolites, within each voxel from the respective intensity spectrum. MRSI is not yet standardized (e.g., using a DICOM format), but there is open source and proprietary software available for pre-processing and metabolite quantification like LCModel (Provencher, 1993), jMRUI (Stefan et al., 2009), and Totally Automatic Robust Quantitation in NMR (TARQUIN) (Reynolds et al., 2006).

Common pre-processing steps provided by the tools are eddy current compensation, offset correction, noise filtering, zero filling, residual water sup-

pression, and phase and base line correction. These steps are executed in the measured time domain or in the frequency domain after performing a Fourier transform. Signal strength in time domain indicates the metabolites' concentration, while the area under the metabolite curve is used to compute the concentration in frequency domain. The computed values are similar when measurements with high signal-to-noise ratio are provided (Vanhamme et al., 2001). For the actual metabolite concentration, mathematical models are used and fitted to the measured data. Different tools use different fitting approaches.

For the pre-processing of our data, we used TARQUIN, which is an open-source GUI-based software for in-vivo spectroscopy data pre-processing and quantification. It operates in time domain and uses a non-negative least-squares method for model fitting to compute the metabolite concentrations. We incorporate TARQUIN in our data preparation step due to free availability, friendly user interface, automatic quantification, support for multi-voxel spectroscopy, and being able to process various file formats and to export the results in various formats (Reynolds et al., 2006; Wilson et al., 2011).

3.3 Data

The data we use in this paper for our application scenario was acquired using an ^{1}H MRSI technology on a 3T Siemens scanner (TR/TE/flip = $1700ms/135ms/90°$). Two MRSI series are taken, each having a $160mm \times 160mm \times 1mm$ field of view. In addition, an MRI volume is captured and registered with the MRSI volume. The MRI volume is $224mm \times 256mm \times 144mm$ with 1 mm slice thickness. We clip the MRI volume to the MRSI volume. The resolution of the MRI volume is much higher such that each MRSI voxel stretches over $10mm \times 10mm \times 12mm$ MRI voxels. The data are courtesy of Miriam Bopp and Christopher Nimsky from the University Hospital Marburg, Germany. We used TARQUIN for MRSI pre-processing and metabolite quantification. The quantification process resulted in 33 metabolites that are listed in Table 1.

After metabolite quantification, we can summarize the available data as two slabs of MRSI voxels with registered MRI voxels. For each MRSI voxel, we have computed concentrations of 33 metabolites, leading to a 33-dimensional feature space, where each point in the feature space represents the chemical composition of one voxel. In addition, we know for each MRSI voxel, which are the matching MRI voxels (single intensity values). In the following, we will propose methods for visualization in image space and feature space, where coordinated interaction within the linked views is used to analyze the MRSI data.

Table 1: Metabolites delivered by a quantification from brain MRSI data using TARQUIN.

No.	Name	No.	Name	No.	Name
1	Ala	12	Lip09	23	PC
2	Asp	13	Lip13a	24	PCr
3	Cr	14	Lip13b	25	Scyllo
4	GABA	15	Lip20	26	Tau
5	GPC	16	MM09	27	TNAA
6	Glc	17	MM12	28	TCho
7	Gln	18	MM14	29	TCr
8	Glth	19	MM17	30	Glx
9	Glu	20	MM20	31	TLM09
10	Ins	21	NAA	32	TLM13
11	Lac	22	NAAG	33	TLM20

3.4 Driving Questions

MRSI data are mainly acquired to analyze tumors. In our application scenario, we look into brain tumors. The first driving question is the tumor classification, i.e., one would want to finds out what type of tumor it is and how its malignancy is rated. To do so, one should analyze the MRSI voxels belonging to the tumor. The question one would like to answer with MRSI measurements is then: What are the chemical compositions of the tumor? The second important question is whether the tissue surrounding the tumor is already infiltrated by the tumor. MRSI allows for a more detailed analysis of the surrounding tissue. The question to be answered is then: How do the chemical composition in areas surrounding the tumor compare to those of the tumor and to those of healthy tissue? Finally, one may want to see whether tumors in other regions occur, i.e., one would ask: Are there other areas that have a similar chemical composition as the tumor?

4 METHODOLOGY

4.1 Image Space Visualization

In order to analyze image regions such as tumors or surrounding tissue, one would need to be able to visually inspect those regions and interactively select them. We support this using a slice-based viewer, which we implemented using the VTK (Schroeder et al., 2006) and ITK (Johnson et al., 2013) libraries. Within the slice, we have two visualization layers. The first layer represents the MRI volume. It pro-

vides the anatomical context for the MRSI data analysis, see Figure 1. A greyscale luminance color map is used, as this is a common standard in MRI visualization. The second layer represents the MRSI volume. Here, individual voxels can be selected interactively by brushing on the image regions. Selected voxels are highlighted by color, see Figure 7(left).

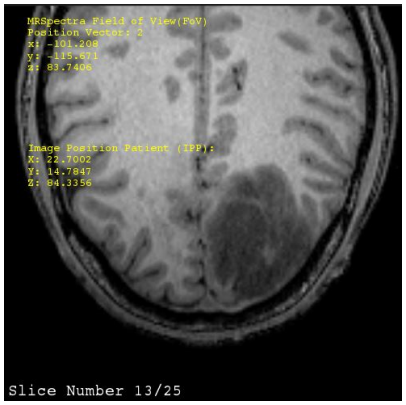

Figure 1: Anatomical context in slice-based image-space MRI visualization.

In addition to manual selection of MRSI voxels based on visual inspection, we also support an automatic image segmentation method. The automatic image segmentation predefines regions for quick selection and region analysis. The automatic segmentation method partitions the MR image. A large range of algorithms exist each having certain advantages and drawbacks. For the data at hand, the best results in our tests were obtained by the Multiplicative Intrinsic Component Optimization (MICO) (Li et al., 2014) segmentation method for auto characterization of various tissues in brain MRI. MICO could handle MRI with low signal-to-noise ratio that is due to magnetic field disturbance and patient movement during the scanning process. MICO performs well in bias field estimation and in discarding intensity inhomogeneity.

If we impose the MR image segmentation result on the MRSI voxels, we have to deal with partial volume effects, as one MRSI voxel corresponds to many MRI voxels, which may be classified differently. We propose to use a visual encoding that conveys the partial volume effect. Instead of simply color-coding the MRSI voxel by the color for the dominant class among the MRI voxels, we render square glyphs that are filled with different colors, where the color portions reflect the percentages how often the MRI voxels are assigned to the respective class. Figure 2 shows a respective slice-based visualization, where the segmentation result is shown in the MRSI layer overlaying the MRI layer. The mixture of colors indicate the uncertainty of the anatomical region segmentation at the MRSI resolution. In the image, we show the glyphs for each of the dominant class, i.e., the five columns (from left to right) show MRSI voxels, where the respective MRI voxels have primarily been classified as background, cerebrospinal fluid (CSF), tumor, gray matter, and white matter, respectively. The colors used for the respective classes are shown above the images. Results are shown for both MRSI slabs.

4.2 Feature Space Visualization

The feature-space visualization problem is that of a multidimensional data visualization problem, where dimensionality is in the range of tens, while the number of data points is in the range of hundreds. Hence, we need an approach that scales sufficiently well in both aspects. Moreover, the tasks require us to observe patterns such as clusters of data points, i.e., sets of data points that are close to each other and different from other data points. Many multidimensional data visualization approaches exist and we refer to a recent survey for their descriptions (Liu et al., 2017). They can be distinguished by the performed data transformations, by the visual encodings, and by their mappings to visual interfaces. Point-based visual encodings scale well in the number of data points and allow for an intuitive detection of clusters. Among them, projection-based dimensionality reduction approaches use data transformations to a visual interface supporting the handling of high dimensionality. Linear projections have the advantage over non-linear projections that the resulting plots can be easily related back to the original dimensions by the means of star coordinates. Thus, we propose to use star-coordinates plots (Kandogan, 2000) to visualize the feature space.

A projection from an n-dimensional feature space to a 2D visual space is obtained by a projection matrix of dimensions $2 \times n$, where each of the n columns represent the tip of the n star-coordinates axes. The default set-up is to locate the tips equidistantly on the unit circle, see Figure 3. The tips can be moved to change the projection matrix, which allows for an interactive multidimensional data space analysis supporting the detection of trends, clusters, outliers, and correlations (Teoh and Ma, 2003).

When assuming an image-space segmentation of the voxels as described in the preceding section, each voxel is assigned to a class. Hence, we are dealing with *labeled* multidimensional data. A common task in the visual analysis of labeled multidimensional data is to find a projected view with a good separation of the classes. This is also of our concern, as we want

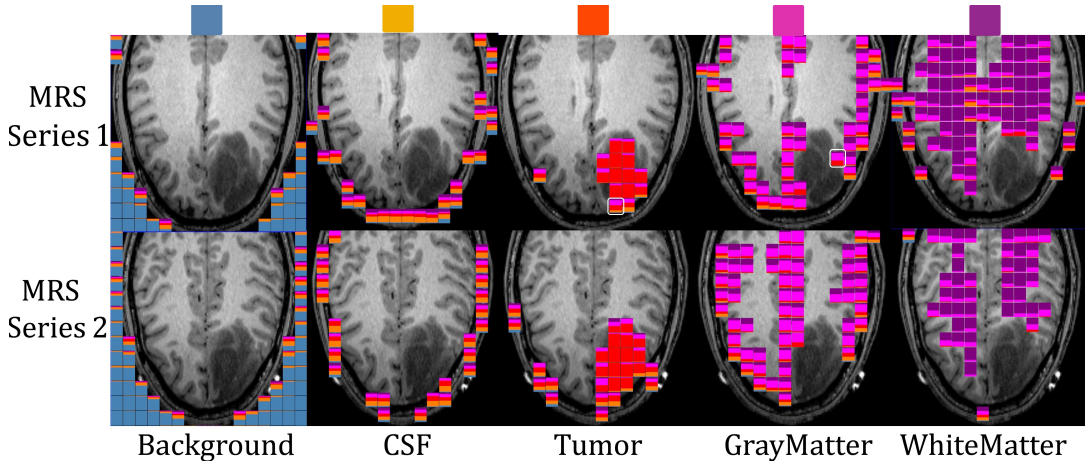

Figure 2: MRI segmentation visualized at MRSI resolution using glyphs with color ratios that reflect class distribution of each voxel.

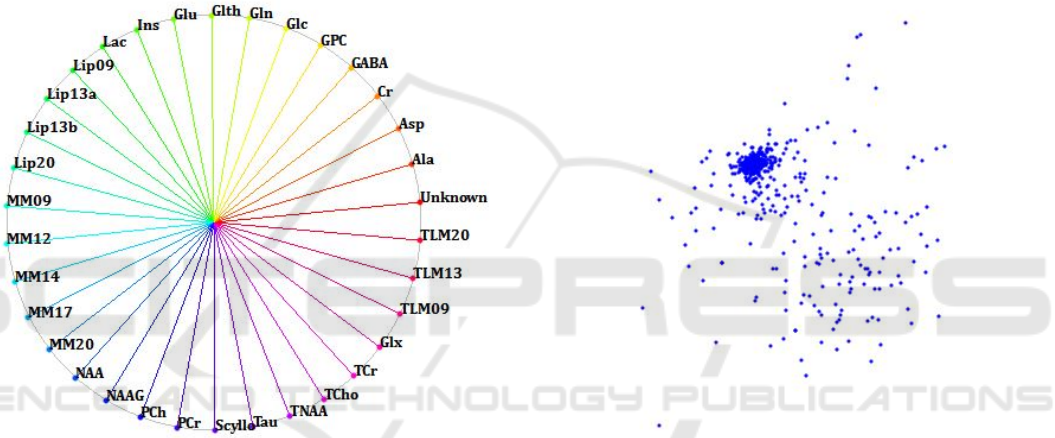

Figure 3: Default star-coordinates configuration (left) and respective projection of 33-dimensional feature space to a 2D visual space. Each dimension represents a metabolite concentration, each point corresponds to a MRSI voxel.

to separate tumor from healthy tissues. Molchanov et al. (Molchanov and Linsen, 2014) proposed an approach for intuitive interactive class separation in linearly projected views. We adopt their idea for our purposes. The idea is to use a control point for each class, e.g., being the classes' medians. Classes can then be separated by dragging control points apart. Since the visual representation is restricted to linear projections, the position to which the control points are dragged can, in general, not be exactly obtained in a linear projection. Molchanov et al. proposed to use a least-squares approach to find the best match to the desired interaction. Using this idea, the user just needs to move the control points of the classes in an intuitive manner, where the number of classes is usually low (five in the case of brain imaging data), instead of interacting with all star-coordinates axes, which becomes tedious for larger number of dimensions (33 in our application scenario). Figure 4 illustrates this by showing the star-coordinates configuration to the left and the linear projection of the labeled multidimensional data to the right, where the control points that can be interactively moved are the ones with a black frame. Since the visualization would be too cluttered when overlaying the two views, we decided to show them in a juxtaposed manner. Further argumentation for juxtaposed views of star coordinates is provided by Molchanov and Linsen (Molchanov and Linsen, 2018).

Since the points in our projected view represent MRSI voxels, while the image segmentation is performed on the higher-resolution MR image, we have partial volume effects as discussed in the previous section. The uncertainty of the labeling result shall also be conveyed in our projected view. For example, if some points are identified as outliers of a cluster, one should be able to reason if the labeling of the respective voxel is uncertain or not. If it was uncertain, the outlyingness may be due to a wrong labeling decision. To visually convey the labeling uncertainty, the

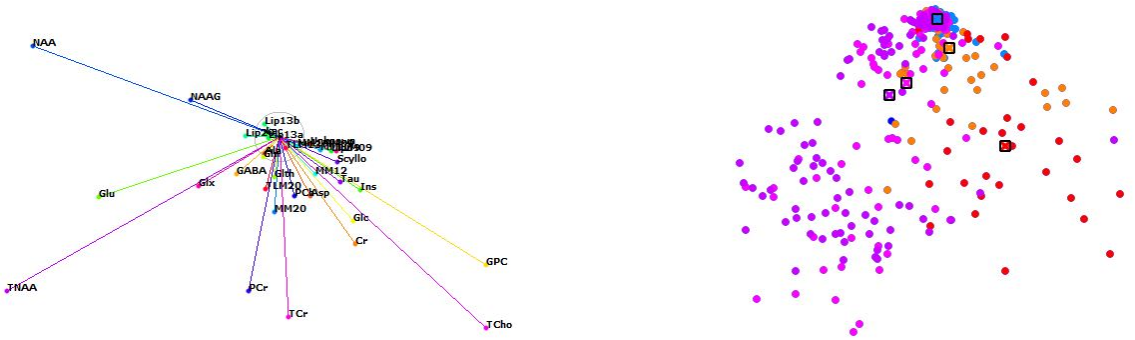

Figure 4: Interactive visual analysis of labeled multidimensional data: star-coordinates configuration (left) and respective projected view (right) obtained by interacting with control points (black frames) of the classes induced by image segmentation.

probabilities of each point belonging to each class label shall be represented. Since the probabilities sum up to one, they are well represented by a pie chart, i.e., each point in the projected view is displayed by a pie chart. Figure 5 shows a projected view with labeling uncertainty conveyed by pie charts.

While interacting with the control points in the projected view, the star-coordinates axes are updated accordingly. Hence, when separating classes in the projected view, one can observe in the star-coordinates view, which axes are mainly responsible for the separation, i.e., which axes are the ones that allow for such a separation. The dimensions that are associated with these axes are subject to further investigations, as one of our tasks was to detect the metabolic compositions of tumors and surrounding tissues. For selected metabolites, we use different statistical plots supporting different analysis steps.

If we are interested in investigating the metabolic compositions of two classes for selected metabolites, we use juxtaposed box plots to show the statistical information of the metabolic concentrations of all voxels that were assigned to each class. The box plots convey the median, the interquartile range, as well as minimum and maximum, see Figure 12.

If we are interested in investigating the interplay of two metabolites, we can analyze their correlation using a 2D scatterplot. Scatterplots are most effective in showing correlation and detecting outliers, see Figure 6(right). Correlation analysis is supported by computing and displaying a regression line. In case more than two metabolites shall be analyzed simultaneously, parallel coordinate plots allow for a good correlation analysis and scale better than scatterplot matrices in the number of dimensions, see Figure 7(right).

4.3 Coordinated Interaction in Image and Feature Space

The interactive visual analysis becomes effective when using coordinated interaction of the image- and feature-space views. In the image space, one can brush the MRSI voxels in the slice-based visualization and group them accordingly. Hence, interactive labeling is supported. Of course, one can also use the labeling that is implied by the MRI segmentation (including the uncertainty visualization). The resulting labeling is transferred to the projected feature-space visualization by using the same colors in both visualizations, see image-space visualization in Figure 2 and respective projected feature-space visualization in Figures 4 and 5. This coordinated interaction allows for investigating whether image segments form coherent clusters in feature space, which dimensions of the feature space allow for a separation of the clusters, and whether there are some outliers. Of course, the same holds true when using the statistical plots as a coordinated feature-space view. Figure 6 shows a brushing in image space and investigation of the selection in a scatterplot visualization of two (previously selected) dimensions, while Figure 7 shows a brushing in image space and investigation of the selection in a parallel coordinate plot of seven (previously selected) dimensions.

The coordinated interaction of image and feature space is bidirectional. Thus, the user may also brush in the feature space, e.g., by selecting a group of points in the feature space that form a cluster, and observe the spatial distribution of the selection in the image space. Again, one can use any of the feature-space visualizations or a combination thereof. For example, in Figure 8, the four voxels to the left are chosen (voxels having low concentration of TNAA relative to TCho and TCr) in the bar chart and the respective highlighting in the image space convey that these

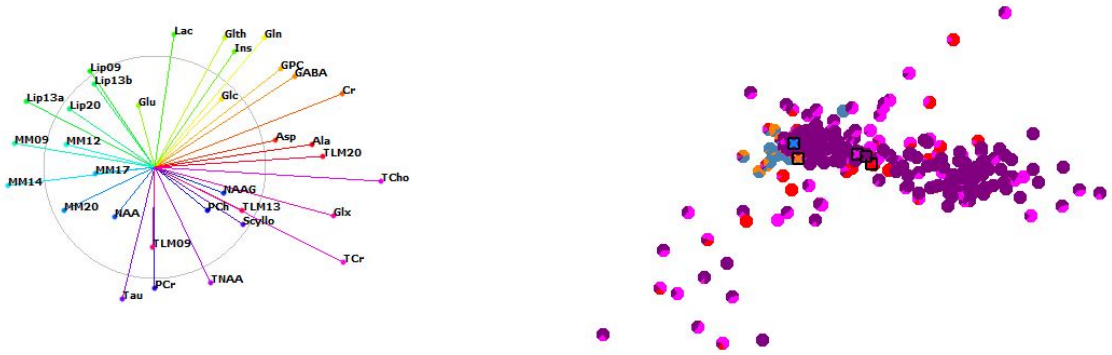

Figure 5: Pie-chart representation of uncertainty for projected points.

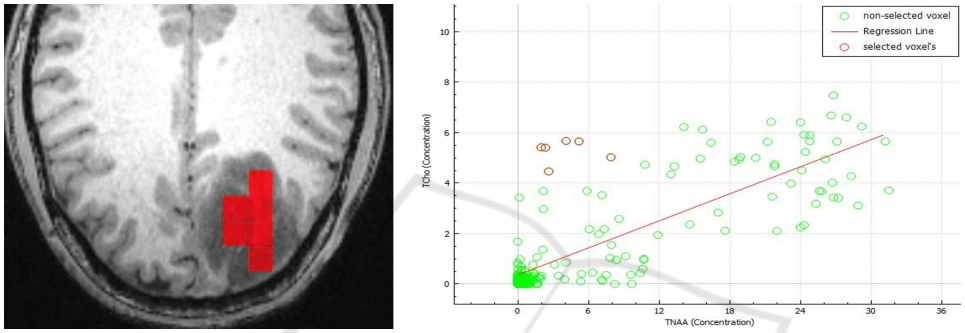

Figure 6: Correlation analysis of the metabolites TCho and TNAA in a scatterplot (right) and detection of outliers. Voxels corresponding to the interactively selected outliers are shown in the image-space visualization (left), which conveys their relation to the tumor.

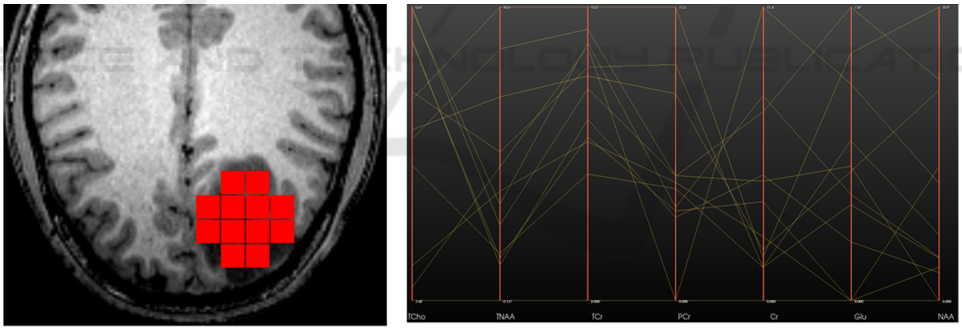

Figure 7: (left) Interactive selection of MRSI voxels in image space and (right) linked view for correlation analysis of seven selected metabolites (TCho, TNAA, TCr, PCr, Cr, Glu, and NAA) using parallel coordinate plot for the selected voxels.

voxels belong to the tumor region.

We also support a heatmap visualization as commonly used when observing the spatial distribution of individual metabolite concentrations. The user may select a single metabolite or a combination of metabolites. Figure 9 shows heatmaps of TCho and TNAA concentrations and In Figure 10, a heatmap of the logarithm of the ratio of TCho/TNAA is shown.

The full potential of our system is reached by using coordinated interaction in both ways simultaneously. For example, one can select the tumor voxels in the image space, can investigate the respective set of points in feature space possibly forming a cluster, detect further points that fall into the cluster, select those further points, and observe their spatial distribution. These newly selected voxels may be voxels surrounding the tumor, in case the tumor has already infiltrated surrounding regions, or may be voxels that form a region elsewhere, in case there is a second tumor.

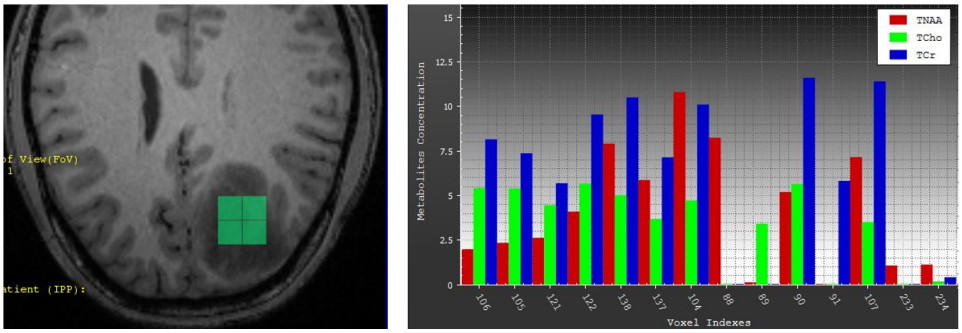

Figure 8: Concentration of three selected metabolites shown side by side in the bar chart (right). Four voxels with low concentration of TNAA relative to TCho and TCr are selected by brushing. The coordinated image-space view (left) highlights the selected voxels, which form the core of the tumor.

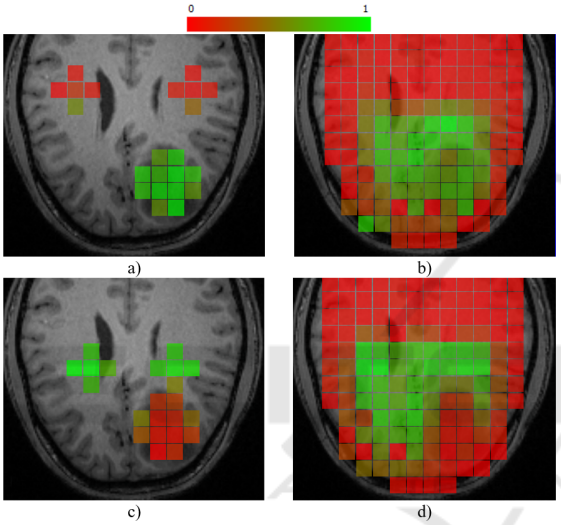

Figure 9: Heatmaps of TCho (top) and TNAA concentration (bottom) for all brain voxels (right) and selected regions of interest (left).

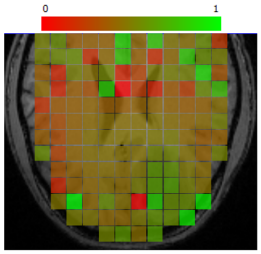

Figure 10: Heatmap of ratio of TCho/TNAA concentrations (in logarithmic scaling).

5 RESULT AND DISCUSSION

In our application scenario, we applied the developed methods of Section 4 to the data acquired from a 26-year old male patient with a brain tumor. MRSI and MRI head scans were obtained as described in Section 3.3. We preprocessed the data as described in Section 3.2. Then, we first investigated the projected feature space using the default star-coordinates layout as in Figure 3. We observe no obvious clusters in the projected space. Thus, we next applied an automatic segmentation of the MR image and imposed the segmentation onto the MRSI voxels using our uncertainty-aware visualization in Figure 2. The extracted segments represent grey matter, white matter, CSF, the tumor, and background. This segmentation implies a labeling as shown in Figure 4, where the colors match with the ones in Figure 2. We applied the interactive technique to separate the classes in feature space using the interaction with the classes' control points. In particular, we applied it to separate the tumor class (red) from the other classes. We observe that certain dimensions of our 33-dimensional feature space are affected strongly by this interactive optimization. Hence, the respective metabolites may be the ones that distinguish tumor from the other segments. In Figure 4, we see that the axes labeled NAA, Glu, TNAA, PCr, TCr, TCho, and GPC are longest. These metabolites are candidates for further investigations. In Figure 6, we select TNAA and TCho concentrations, which had the longest axes in Figure 4 and plot all voxels in a scatterplot. We select a group of outliers (red) with high TCho and low TNAA concentrations and observe that these form the core of the tumor region. Please note that by just looking at one of the two metabolites the voxels would have not been outliers. In Figure 9, we use heatmaps to investigate the spatial distributions of TCho and TNAA concentrations, respectively. When looking at the entire brain region (right image column), it is hard to detect structures. However, when selecting regions of interest such as a white matter region and the tumor region, we can spot the differences (left image column). The two columns apply the red-to-green color map to the minimum-to-maximum range of selected voxels only. In Figure 10, a heatmap of the logarithm of the ratio of the two metabolites is shown.

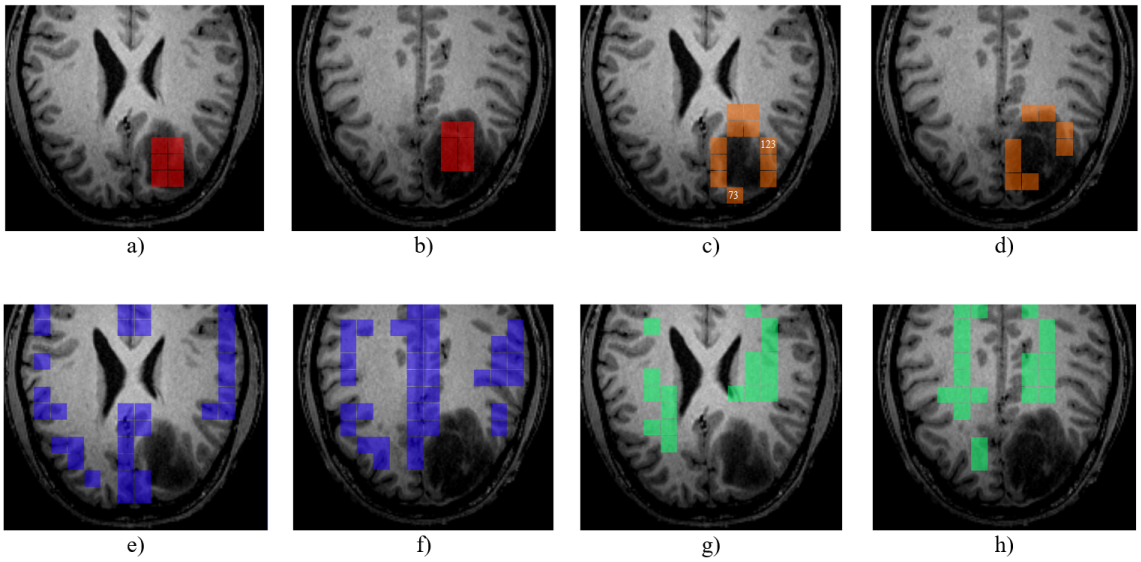

Figure 11: Interactive labeling of voxels within both MRSI slabs representing tumor (red), voxels in the vicinity of the tumor (orange), grey matter (blue), and white matter (green).

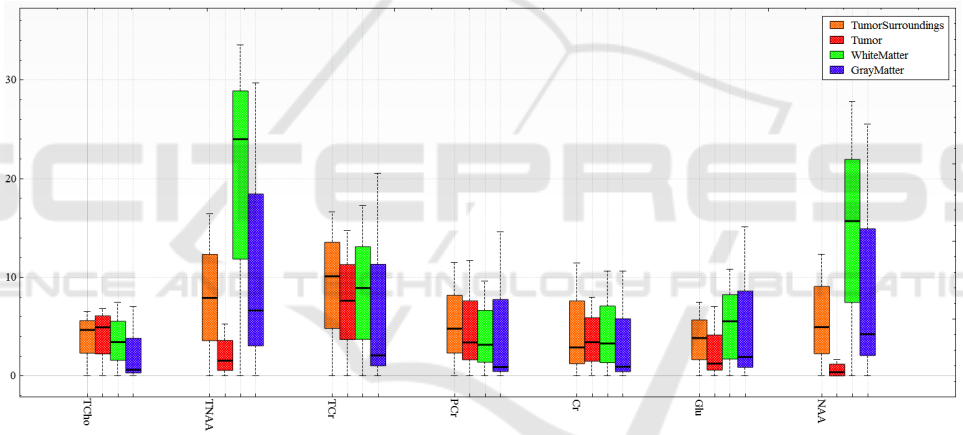

Figure 12: Juxtaposed box plots to compare class statistics for selected classes (voxels labeled as white matter, grey matter, tumor, and surrounding tumor) and selected metabolites.

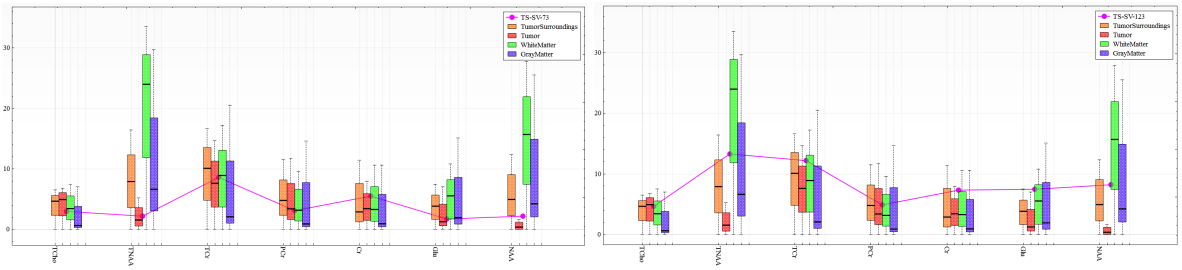

Figure 13: Metabolite concentration of voxels 73 (left) and 123 (right) in vicinity of the tumor selected in Figure 11 in comparison to the concentrations of the tumor, white matter, and grey matter classes.

When investigating the projected feature space in Figure 4, we observe that the classes are actually not well separated. We further investigate the pie chart visualization of the labeling result, see Figure 5 to observe quite a few voxels that are rather uncertain with respect to the automatic labeling, which may be due to the partial volume effect. Also, when selecting all voxels that contain parts of the tumor as

in Figure 7, the distribution of values in the parallel coordinate plot appear rather diverse. Thus, we decided to manually select regions of low-uncertainty voxels and create new labels, see Figure 11. Figure 12 shows the juxtaposed box plots of the four selected voxel groups. We observe that the tumor region (red) is quite different from the other regions in NAA and TNAA concentrations, where TNAA is the sum of NAA and NAAG. We also observe that the voxels surrounding the tumor (orange) behave like gray/white matter voxels rather than tumor voxels for these metabolites, which makes us believe that these voxels are not yet infiltrated by the tumor. However, this may not be true for individual voxels of that group. Since the tumor size and shape of its boundary is important for diagnosis and treatment, we selected individual voxels at the border of the tumor and investigated the metabolite concentrations and individually compared their metabolite concentrations to those of tumor as well as white and grey matter. Figure 13 shows such a comparison for the voxels labeled 73 and 123 in Figure 11. These two voxels showed high uncertainties in the automatic segmentation result. We can observe in Figure 13 that voxel 73 matches well the tumor class, while voxel 123 does not. Hence, we conclude that the tumor may already have infiltrated the area at voxel 73, while it may have not yet done so for voxel 123.

We invited two MRSI experts with many years of experience of acquiring and analyzing MRSI data to our lab to show them our tool on a large multi-touch display. The visual encodings were mostly intuitive to them and they were quickly able to suggest interactive analysis steps themselves. Only the projected view needed some explanations, but the intuitive interaction with the control points made them adopt the concept quickly. They also quickly brought in their expertise into the analytical workflow by excluding some metabolites that they knew would not be important for the given tasks such as lipids and macromolecules and by interpreting correctly combinations such as TNAA being a combination of NAA and NAAG. In the session, we jointly looked into the metabolic composition of voxels in the vicinity of the tumors as documented above. To test the reproducibility of the analysis, it would be desirable to test our tool with a large number of experts on a large number of data sets. We hope that we can conduct such a study in future work, but acknowledge that it will be challenging to recruit a large number of experts.

6 CONCLUSIONS AND FUTURE WORK

We presented a comprehensive tool for the analysis of MRSI data and applied it to brain tumor investigations. Using coordinated views of image and feature space visualizations, effective analysis steps can be performed that allow for a comprehensive investigation of all data facets. We showed how relevant metabolites can be identified and how image regions can be detected and compared. In this paper, we focused on the analysis of a tumor region. Future directions include the analysis of different tumor types, for which MRSI information from many patient shall be combined, eventually leading to a cohort analyses.

ACKNOWLEDGEMENTS

This work was supported in part by DFG grant MO 3050/2-1.

REFERENCES

Board, PDQ Adult Treatment Editorial (2018). Adult central nervous system tumors treatment PDQ®. In *PDQ Cancer Information Summaries*. National Cancer Institute (US).

Burnet, N. G., Thomas, S. J., Burton, K. E., and Jefferies, S. J. (2004). Defining the tumour and target volumes for radiotherapy. *Cancer Imaging*, 4(2):153–161.

Crane, J. C., Olson, M. P., and Nelson, S. J. (2013). SIVIC: open-source, standards-based software for DICOM MR spectroscopy workflows. *Journal of Biomedical Imaging*, 2013:12.

Feng, D., Kwock, L., Lee, Y., and Taylor II, R. M. (2010). Linked exploratory visualizations for uncertain MR spectroscopy data. In Park, J., Hao, M. C., Wong, P. C., and Chen, C., editors, *Visualization and Data Analysis 2010*. SPIE.

Gujar, S. K., Maheshwari, S., Björkman-Burtscher, I., and Sundgren, P. C. (2005). Magnetic resonance spectroscopy. *Journal of neuro-ophthalmology*, 25(3):217–226.

Johnson, H. J., McCormick, M., Ibáñez, L., and Consortium, T. I. S. (2013). *The ITK Software Guide*. Kitware, Inc., third edition.

Kandogan, E. (2000). Star coordinates: A multi-dimensional visualization technique with uniform treatment of dimensions. In *In Proceedings of the IEEE Information Visualization Symposium, Late Breaking Hot Topics*, pages 9–12.

Li, C., Gore, J. C., and Davatzikos, C. (2014). Multiplicative Intrinsic Component Optimization (MICO) for MRI bias field estimation and tissue segmentation. *Magnetic resonance imaging*, 32(7):913–923.

Liu, S., Maljovec, D., Wang, B., Bremer, P., and Pascucci, V. (2017). Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization & Computer Graphics*, 23(3):1249–1268.

Matkovic, K., Freiler, W., Gracanin, D., and Hauser, H. (2008). ComVis: A coordinated multiple views system for prototyping new visualization technology. In *2008 12th International Conference Information Visualisation*, pages 215–220. IEEE.

Maudsley, A., Darkazanli, A., Alger, J., Hall, L., Schuff, N., Studholme, C., Yu, Y., Ebel, A., Frew, A., Goldgof, D., et al. (2006). Comprehensive processing, display and analysis for in vivo MR spectroscopic imaging. *NMR in Biomedicine*, 19(4):492–503.

McKnight, T. R., Noworolski, S. M., Vigneron, D. B., and Nelson, S. J. (2001). An automated technique for the quantitative assessment of 3D-MRSI data from patients with glioma. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 13(2):167–177.

Molchanov, V. and Linsen, L. (2014). Interactive design of multidimensional data projection layout. In N. Elmqvist, M. Hlawitschka, and J. Kennedy, editors, *EuroVis - Short Papers*. The Eurographics Association.

Molchanov, V. and Linsen, L. (2018). Shape-preserving star coordinates. *IEEE Transactions on Visualization & Computer Graphics*, pre-print.

Nunes, M., Laruelo, A., Ken, S., Laprie, A., and Bühler, K. (2014a). A survey on visualizing magnetic resonance spectroscopy data. In *Proceedings of the 4th Eurographics Workshop on Visual Computing for Biology and Medicine*, pages 21–30. Eurographics Association.

Nunes, M., Rowland, B., Schlachter, M., Ken, S., Matkovic, K., Laprie, A., and Bühler, K. (2014b). An integrated visual analysis system for fusing MR spectroscopy and multi-modal radiology imaging. In *Visual Analytics Science and Technology (VAST), 2014 IEEE Conference on*, pages 53–62. IEEE.

Provencher, S. W. (1993). Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magnetic resonance in medicine*, 30(6):672–679.

Raschke, F., Jones, T., Barrick, T., and Howe, F. (2014). Delineation of gliomas using radial metabolite indexing. *NMR in Biomedicine*, 27(9):1053–1062.

Reynolds, G., Wilson, M., Peet, A., and Arvanitis, T. (2006). An algorithm for the automated quantitation of metabolites in in vitro NMR signals. *Magnetic resonance in medicine*, 56(6):1211–1219.

Rowland, B., Deviers, A., Ken, S., Laruelo, A., Ferrand, R., Simon, L., and Laprie, A. (2013). Beyond the metabolic map: an alternative perspective on MRSI data. *ESMRMB 2013, 30th Annual Scientific Meeting*, page 270.

Scheenen, T. W., Heerschap, A., and Klomp, D. W. (2008). Towards ^{1}H-MRSI of the human brain at 7T with slice-selective adiabatic refocusing pulses. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 21(1-2):95–101.

Schroeder, W., Martin, K., Lorensen, B., Avila, Sobierajski, L., Avila, R., and Law, C. (2006). *The Visualization Toolkit*. Kitware, Inc., fourth edition.

Stefan, D., Di Cesare, F., Andrasescu, A., Popa, E., Lazariev, A., Vescovo, E., Strbak, O., Williams, S., Starcuk, Z., Cabanas, M., et al. (2009). Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package. *Measurement Science and Technology*, 20(10):104035.

Teoh, S. T. and Ma, K.-L. (2003). StarClass: Interactive visual classification using star coordinates. In *SDM*, pages 178–185. SIAM.

Vanhamme, L., Sundin, T., Hecke, P. V., and Huffel, S. V. (2001). MR spectroscopy quantitation: a review of time-domain methods. *NMR in Biomedicine*, 14(4):233–246.

Wilson, M., Reynolds, G., Kauppinen, R. A., Arvanitis, T. N., and Peet, A. C. (2011). A constrained least-squares approach to the automated quantitation of in vivo ^{1}H magnetic resonance spectroscopy data. *Magnetic resonance in medicine*, 65(1):1–12.

Wolf, I., Vetter, M., Wegner, I., Nolden, M., Bottger, T., Hastenteufel, M., Schobinger, M., Kunert, T., and Meinzer, H.-P. (2004). The Medical Imaging interaction ToolKit MITK: A toolkit facilitating the creation of interactive software by extending VTK and ITK. volume 5367, pages 5367 – 5367 – 12.