Multi-View Visual Analysis of Chemical Shifts in Magnetic Resonance Spectroscopy

Tim Markwardt Antonios Raptakis Nicola M. Drüeke Davy Göbel Steve Dübel
Lars Nonnemann*

Fraunhofer Institute for Computer Graphics Research (IGD)

ABSTRACT

This paper presents a novel multi-view visualization system for analyzing biochemical information in 31P Magnetic Resonance Spectroscopy (MRS) data. We combine Savitzky-Golay denoising and automated peak detection with coordinated views to support metabolite identification, quality assessment, and targeted comparative analysis. The system includes: Overview (rapid triage), Horizon Graphs (compact comparison vs user-selected reference), and Single Spectrum/Peak Analysis (detailed peak-area and ratio inspection).

Index Terms: Interactive Visualization, Magnetic Resonance Spectroscopy

1 IDENTIFYING THE PROBLEM

Neurological conditions are a leading cause of illness and disability worldwide according to the World Health Organization (WHO), which necessitates innovative approaches for their early detection and diagnosis. Magnetic Resonance Spectroscopy (MRS) offers the capability to non-invasively assess biochemical changes in the brain, thereby providing a valuable contribution to medical research. Here, interactive data visualization plays a crucial role in the analysis and interpretation of such complex medical data. Physicians typically focus on a small set of metabolite peaks rather than entire spectra. Therefore, visualization designs should emphasize these peaks and allow direct comparison. Existing MRS visualization approaches typically rely on full-spectrum line plots with annotated peaks and fitted-model overlays (e.g., jMRUI [9], LCModel[5]) and, for spectroscopic imaging, spatial heatmaps/topographic maps of metabolite distributions. Furthermore, research prototypes add small-multiples and linked/coordinated views [4], [2], [8], but clinical tools often remain focused on static full-spectrum displays and batch quantification rather than interactive, metabolite-focused comparison workflows [9], [5].

To address these needs, we designed a system to de-emphasize complete full-spectrum readouts for routine tasks and instead make targeted metabolite comparison a primary design goal. By adapting coordinated views and horizon graphs to address the challenges of MRS data interpretation, we intend to contribute to the development of effective visualization techniques that benefit both clinical researchers and medical practitioners.

2 REDESIGNING THE 31P-MRS VISUALIZATION

To extract features, such as metabolite peaks, from an inherent noisy signal produced by 31P magnetic resonance spectroscopy (MRS), we first applied a smoothing filter to the spectral data, specifically utilizing the Savitzky-Golay filter as introduced by Savitzky and Golay [7] and implemented in the SciPy Python library [10]. We selected the filter due to its superior ability to preserve peak shape and position while effectively reducing noise—critical characteristics for accurate metabolite identification in spectroscopy data. This approach allowed us to increase the precision of the data without distorting the underlying signal trend. Figure 1 shows the original signal intensity (blue) alongside the denoised line (orange). To facilitate the identification of metabolites, we detected peaks on the denoised signal, as indicated by the red circles. However, to find the peak and respective chemical shift on the original, unfiltered signal, each detected region $\pm$ 1ppm around the red peak (shown in green), was defined as an arbitary area to identify the maximum intensity on the unfiltered signal and project it to the x-axis, as is pointed out by the black arrows in Fig. 1. This per-peak windowing focuses analysis on metabolite-specific regions; we emphasize peak-area comparisons (rather than peak height), which better reflect metabolite concentrations and support assessment of pathological changes [3]. To provide a comprehensive and user-friendly analysis environment, we propose a three-part solution, structured as interactive tabs:

Overview Tab

The Overview tab (Figure 2) displays multiple selected spectra as line plots with detected metabolite peaks marked by dotted vertical lines, providing a concise dataset summary while reducing visual clutter. It supports rapid triage — quick quality checks, outlier detection, and comparison of peak presence, position, and relative intensity across subjects or time points. By emphasizing detected metabolite peaks, the tab discourages routine full-spectrum comparisons for typical clinical tasks. Users can quickly select spectra of interest for deeper inspection in the Single Spectrum/Peak Analysis Tab.

Horizon Graph Tab

Horizon graphs present compact, banded, color-coded representations of spectra aligned to a user-selected reference, enabling rapid cohort and longitudinal comparisons while saving vertical space [6]. In our implementation, each spectrum is transformed into a horizon graph and aligned to a user-chosen reference spectrum, which serves as the standard for comparison. This design choice offers several advantages: it enables clinicians to select high-quality reference spectra free from artifacts or noise, ensuring meaningful comparisons; it supports targeted clinical questions by allowing comparison against specific patient baselines (e.g., pre-treatment spectra) or established control cases and it accommodates the inherent variability in MRS data where population averages might obscure clinically relevant individual differences due to factors like age or pathology. However, this approach also has limitations, as it requires domain expertise to select appropriate references and may introduce bias if unsuitable references are chosen. This design specifically enables key user tasks including longitudinal patient monitoring by comparing follow-up spectra against individual baselines, quality control assessment by comparing against known high-quality reference spectra, pathology detection by comparing

*e-mail: lars.nonnemann@igd-r.fraunhofer.de

patient spectra against healthy control references, and treatment evaluation by using pre-intervention spectra as references for post-treatment comparison. Furthermore, by comparing each spectrum to this chosen reference, users can efficiently detect global shifts in spectral positioning or changes in overall intensity patterns.

Single Spectrum/Peak Analysis Tab

The Single Spectrum/Peak Analysis (cf. Figure 3) tab provides a comprehensive, multi-view approach for the detailed examination of individual spectra. This interface consists of three coordinated visualizations that complement each other to enable a thorough analysis of metabolite peaks and their relationships. A key aspect in the clinical interpretation of 31P-MRS data is the quantification of the area under each peak, as this measure is more indicative of metabolite concentration and potential pathological changes than peak height alone [1]. While our current implementation focuses on visual representation of peak areas to make these relationships and ratios immediately accessible, future work will incorporate sophisticated quantification methods for precise area calculation. At the top, a horizontal stacked bar chart displays the proportional contribution of each detected peak to the total area (marked in colored regions Figure 3(a)). This visualization, with colors corresponding to the peaks in the spectrum view below, allows for an intuitive comparison of relative metabolite abundance. Directly below (cf. Figure 3(b)), the main spectrum line chart presents the selected spectrum with peaks indicated by a colored background. This view enables users to visually assess the shape and prominence of each peak, supporting the identification of expected metabolites as well as the evaluation of spectral quality. For further analysis, users can select individual peak areas for closer inspection (Figure 3(c)). When a peak is selected, the background of the bottom-most spectrum plot is highlighted in the corresponding color, reinforcing the connection between the quantitative and visual representations.

All visualizations provide standard interactivity such as zooming, panning, and selection. This allows users to closely examine regions of interest and facilitates a more detailed exploration of spectral features. Together, these three coordinated views enable users to efficiently analyze the composition of individual spectra, compare metabolite ratios, and investigate the significance of specific peaks in the context of clinical or research questions.

3 Discussion and Future Work

Our multi-view approach integrates denoising and coordinated visualizations to improve metabolite identification, quality assessment, and targeted comparisons compared to standard full-spectrum viewers. We prioritize metabolite-focused workflows and plan predefined metabolite sets and region filters co-developed with domain experts to reduce overload and support routine clinical tasks. The visualization solution we have developed emphasizes visual exploration and interactive analysis, allowing users to gain insights directly from the data. By explicitly visualizing peak areas under the curve and their ratios in the Single Spectrum/Peak Analysis view, our solution provides relevant information for clinical and research interpretation than standard spectral plots.

Together with a medical physicist from the University Medical Centre Rostock, we examined our solution in a first interview. It was noticed that the comparison of spectra across subjects is not common in clinical practice, occurring at most occasionally in research contexts, primarily due to variations in peak positions caused by pH shifts or magnetic field inhomogeneity. While our Horizon Graph visualization (cf. Figure 4) supports such comparisons, its utility is currently limited by the lack of appropriate filtering options and insufficient focus on regions with high information density. The Detail Plot (Figure 3(c)), while potentially useful for examining complex peak structures, may have limited application in routine clinical analysis. The physicist indicated, detailed analysis at this level is typically reserved for special cases, such as closely positioned or multiple peaks, and is more common in NMR spectroscopy than in clinical MRS applications.

For future work, a significant enhancement to our visualization would be the implementation of predefined sets of metabolites for comparative analysis. Rather than comparing entire spectra, focusing on specific metabolites of interest would align better with clinical practice. Such predefined sets would facilitate more targeted analyses and could be developed in collaboration with domain experts to ensure clinical relevance. The horizon graph visualization, while powerful for comparative analysis, requires additional filtering capabilities to improve clarity and user control. Future implementations could include options to filter by (1) Metabolite of interest – allowing users to focus only on specific peaks Signal quality thresholds – enabling the exclusion of noisy or low-quality spectra (2) Signal quality thresholds – enabling the exclusion of noisy or low-quality spectra (3) Deviation magnitude – highlighting only significant differences from reference spectra (4) Peak alignment quality – excluding spectra with significant peak position drift due to pH shifts or field inhomogeneity, directly addressing the major limitation identified in cross-subject spectral comparisons These enhancements would address the feedback received from our domain expert and make the horizon graph more valuable for both clinical and research applications. A critical next step is a systematic comparison of our visualization-based approach with established quantification methods used in tools like jMRUI and LCModel. While our current implementation focuses on visual exploration, integrating more sophisticated quantification algorithms would enhance the analytical capabilities of the system.

References

[1] N. S. Al-Akl et al. Untargeted metabolomic profiling reveals differentially expressed serum metabolites and pathways in type 2 diabetes patients with and without cognitive decline: A cross-sectional study. *International Journal of Molecular Sciences*, 25(4), 2024. doi: 10.3390/ijms25042247 2

[2] J. C. Crane et al. SIVIC: Open-Source, Standards-Based Software for DICOM MR Spectroscopy Workflows. *International Journal of Biomedical Imaging*, pp. 1–12, 2013. doi: 10.1155/2013/169526 1

[3] I. B. Ip and H. Bridge. Investigating the neurochemistry of the human visual system using magnetic resonance spectroscopy. *Brain Structure and Function*, 227(4):1491–1505, 2022. doi: 10.1007/s00429-021-02273-0 1

[4] A. A. Maudsley et al. Comprehensive processing, display and analysis for in vivo MR spectroscopic imaging. *NMR in Biomedicine*, (4):492–503, 2006. doi: 10.1002/nbm.1025 1

[5] S. Provencher. *LCModel1 & LCMgui User's Manual*, 2024. 1

[6] H. Reijner et al. The development of the horizon graph. In *Electronic Proceedings of the VisWeek Workshop From Theory to Practice: Design, Vision and Visualization*, vol. 3, 2008. 1

[7] A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. doi: 10.1021/ac60214a047 1

[8] B. J. Soher et al. Vespa: Integrated applications for RF pulse design, spectral simulation and MRS data analysis. *Magnetic Resonance in Medicine*, 90(3):823–838, 2023. doi: 10.1002/mrm.29686 1

[9] D. Stefan et al. Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package. *Measurement Science and Technology*, 20(10), 2009. doi: 10.1088/0957-0233/20/10/104035 1

[10] P. Virtanen et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2 1

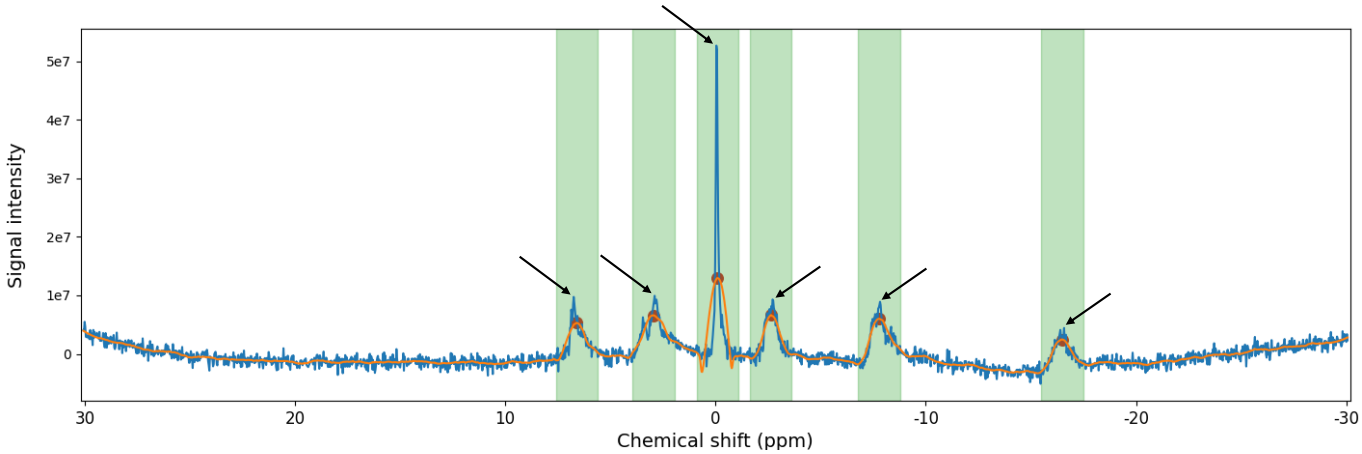

Figure 1: The plot of signal intensity of the chemical shift 10. The blue line indicates the original chemical shift and the orange the denoised line. The red circles correspond to the local peaks. The arrows point out the peaks as they have been found after defining a region, shown with green color, around the peaks of the denoised line.

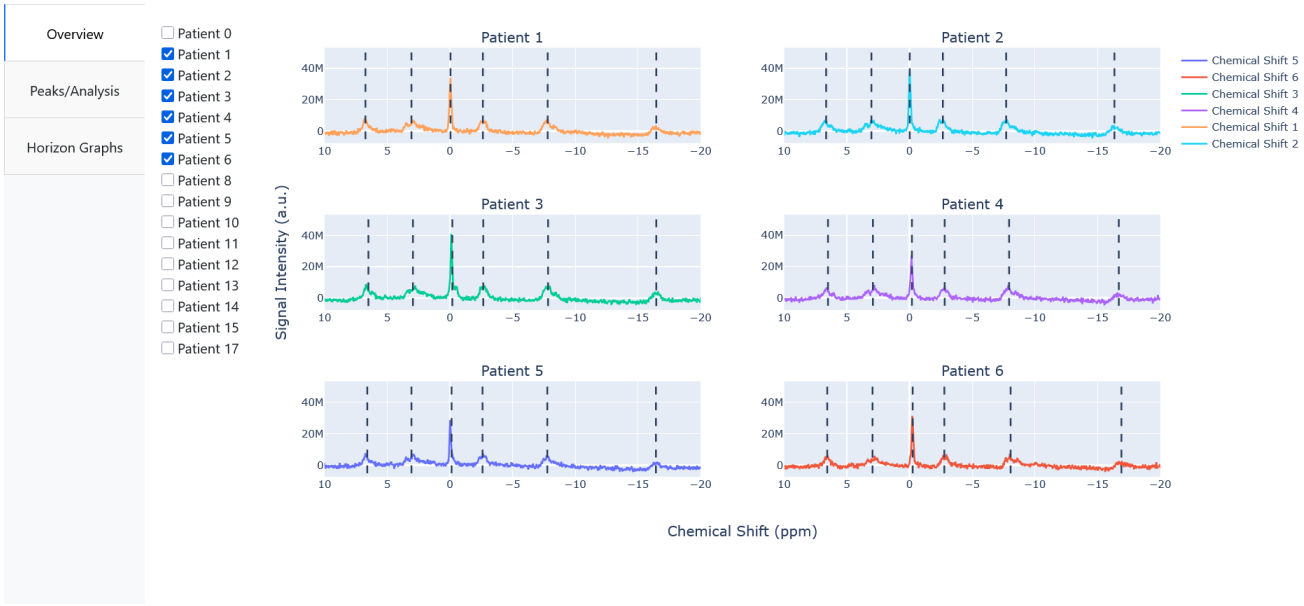

Figure 2: Overview tab displaying spectra from six user chosen patients. Each spectrum is plotted as a non-smoothed line graph with detected metabolite peaks marked by dotted vertical lines.

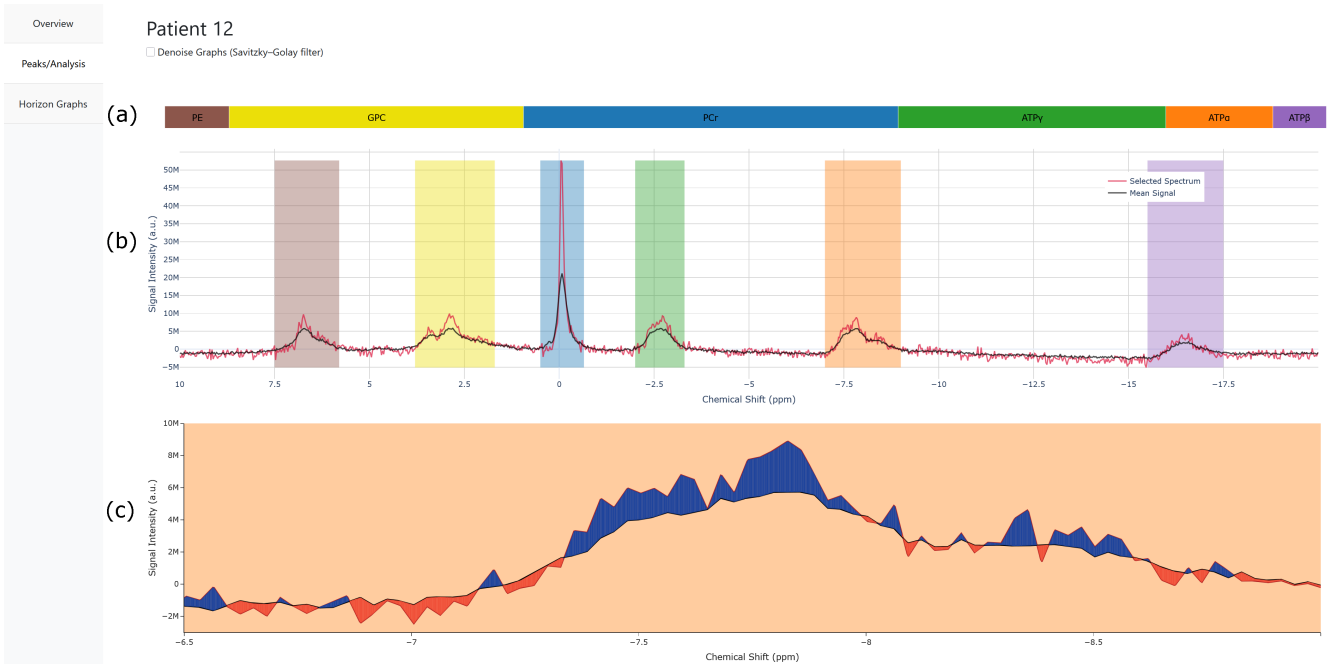

Figure 3: Single Graph/Peak Analysis tab for Patient 12. The top horizontal bar graph **(a)** shows the proportional contribution of the area under the curve (AUC) of each peak to the total AUC of all peaks combined; additionally, the metabolite whose position is closest to each detected peak is labeled in the horizontal bar chart. Peaks are color-coded across all three graphs for consistency. The middle graph **(b)** displays the detected peaks (i.e., metabolites) in the selected spectrum, with the patient spectrum shown in red and a reference graph (black line) overlaid for comparison. The reference graph represents the mean signal of all 17 spectra, chosen here as a placeholder. Users can toggle between the smoothed (Savitzky-Golay filtered) and non-smoothed versions of the spectrum using a checkbox. The bottom graph **(c)** provides a close-up view of a selected peak, highlighting the difference between the patient spectrum and the reference graph by shading the area between them. This detailed view enables precise peak-wise comparison and analysis.

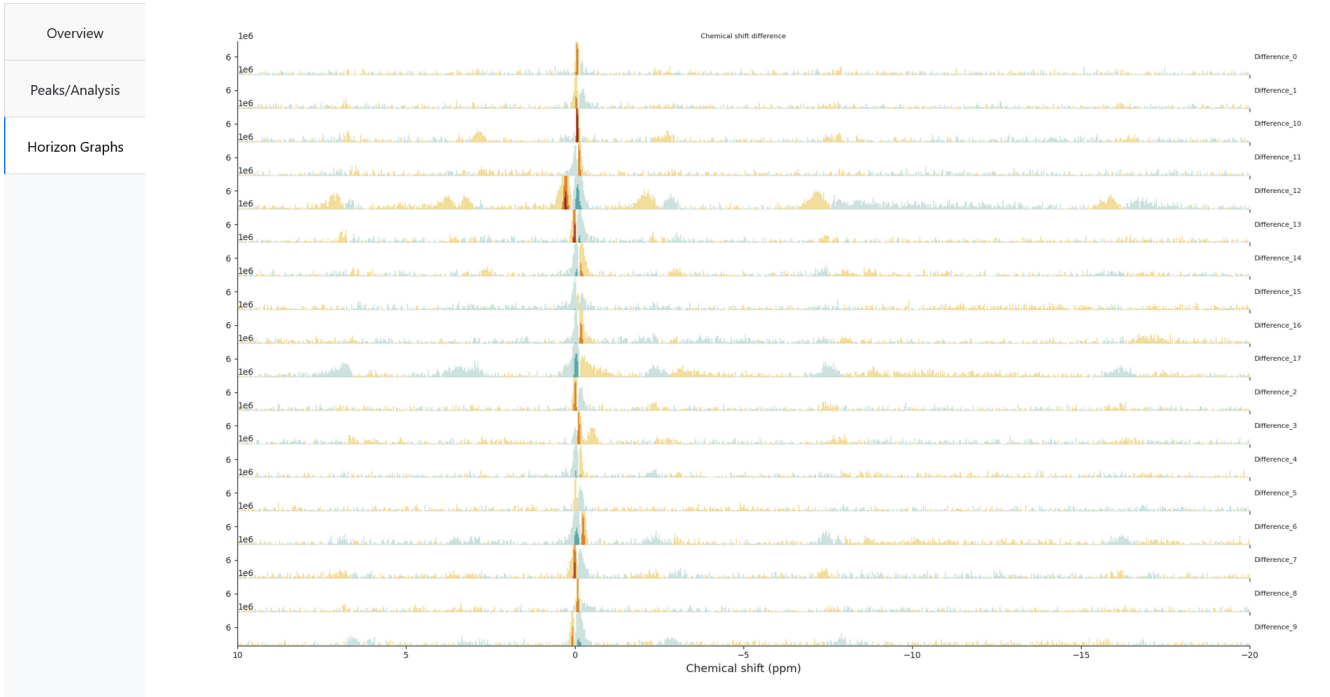

Figure 4: Horizon Graph tab displaying compact, banded, color-coded representations of spectra from 18 patients. Each row corresponds to a patient's spectrum, with deviations from the reference highlighted by color intensity.