MSI-VISUAL: New visualization methods for mass spectrometry imaging and tools for interactive mapping and exploration of m/z values

Jacob Gildenblat[1,2] and Jens Pahnke[1,2,3,4*]

[1]Translational Neurodegeneration Research and Neuropathology Lab, Department of Clinical Medicine (KlinMed), The Medical Faculty, University of Oslo (UiO) and Section of Neuropathology Research, Department of Pathology, Division of Laboratory Medicine (KLM), Oslo University Hospital (OUS), Sognsvannsveien 20, Oslo, NO-0372, Norway, www.pahnkelab.eu.
[2]Institute of Nutritional Medicine (INUM), University of Lübeck (UzL) and University Medical Center Schleswig-Holstein (UKSH), Ratzeburger Allee 160, Lübeck, D-23538, Germany.
[3]Department of Neuromedicine and Neuroscience, The Faculty of Medicine and Life Sciences, University of Latvia (LU), Jelgavas iela 3, Riga, LV-1004, Latvia.
[4]Department of Neurobiology, School of Neurobiology, Biochemistry and Biophysics, The Georg S. Wise Faculty of Life Sciences, University of Tel Aviv (TAU), G. S. Wise Street, Ramat Aviv, IL-6997801, Israel.

*Corresponding author(s). E-mail(s): jens.pahnke@gmail.com;
Contributing authors: jacob.gildenblat@gmail.com;

Abstract

Mass Spectrometry Imaging (MSI) generates high-dimensional molecular data, offering detailed biochemical insights. However, existing MSI software relies primarily on manual pixel-wise m/z selection, limiting the efficiency and depth of data exploration. There is a critical need for computationally advanced, interactive, and automated visualization tools tailored for MSI analysis.
We introduce *MSI-VISUAL*, an open-source framework for fast, interactive, and intuitive MSI data exploration. It enables automatic and manual region selection, statistical comparisons, and real-time visual verification.

MSI-VISUAL integrates established dimensionality reduction techniques and introduces four novel visualization methods for MSI. *Saliency (SALO)* and *Spearman (Spear) Optimizations* directly optimize for global structure preservation and robustness across different distance metrics, addressing key challenges in MSI visualization. *TOP3* and *Percentile Ratio (PR3D)* provide memory-efficient and computationally fast visualizations, enabling efficient analysis of large-scale MSI datasets.

Through comprehensive benchmarking across diverse MSI images, we demonstrate that our proposed visualizations outperform existing techniques, revealing hidden anatomical details with high global structure preservation. We further showcase *MSI-VISUA*L's capability for rapid visual analysis of pathological structures, efficiently mapping m/z distributions for brain atlas sub-regions without MRI registration. Additionally, we map specific m/zs for $A\beta$-plaques, brain tumors, kidney regions, and more, uncovering novel biological insights—such as shared lipid compositions between spinal cord neurons and brain $A\beta$-plaques.

By enabling efficient, interactive, and biologically meaningful MSI data interpretation, *MSI-VISUAL* serves as a powerful resource for both research and clinical applications.

Keywords

mass spectrometry imaging, medical imaging visualization, dimensionality reduction, segmentation, saliency optimization, global structure preservation, percentile ratio, clustering, PCA, NMF, Kmeans, PaCMAP, UMAP, t-SNE, SALO, Spear, PR3D, TOP3, Alzheimer, β-amyloid, brain, kidney, spinal cord

Abbreviations

DR – dimensionality reduction
$A\beta$ – amyloid-β
AD – Alzheimer's disease
MSI – imaging mass spectrometry
NMF3D – 3-dimensional non-negative matrix segmentation
PaCMAP – Pairwise Controlled Manifold Approximation
PCA – principal component analysis
PR3D – 3-dimensional percentile ratio segmentation
ROI – region of interest
SALO – saliency optimization
Spear – Spearman correlation optimization
TIC – total ion current
TOP3 – top 3 m/z segmentation
TriMap – triplet constraints mapping
t-SNE – t-distributed stochastic neighbor embedding
UMAP – uniform manifold approximation

1 Introduction

1.1 Mass spectrometry imaging

Mass spectrometry imaging (MSI) can generate datasets for visualizations that reveal biological and functional structures based on metabolites, lipids, and proteins. It can therefore provide valuable information not accessible through histopathology or other imaging modalities [1–4]. MS images are typically represented as high-dimensional images of shape Rows x Columns x N, where N is the number of molecular measurements per point (also referred to as MS spectra). Each pixel in the image has an associated mass/charge (m/z) spectrum with N values. If an m/z measurement does not align perfectly with a predefined bin, its intensity is typically summed into the closest bin. Mapping from N to 1 or 3 dimensions is needed to visualize these high-dimensional images for humans.

Next, we give a short overview of existing software for MSI and previous computational visualization methods applied for MSI, before describing the *MSI-VISUAL* framework and the proposed new visualizations.

1.2 Previous software for MSI analysis

The *Cardinal* [5] and *SCiLS Lab* (Bruker Daltonics, Bremen, Germany) are popular MSI analysis software. *Cardinal* focuses on computational features in the R programming language, while *SCiLS Lab* is a user interface for MS images scanned with Bruker hardware. Both *Cardinal* and *SCiLS Lab* allow investigating m/z profiles for clicked data points. These may have thousands (or tens of thousands) of m/z values, making it impractical to explore, biasing users to check m/z values with the highest intensities that are easier to notice. The recently proposed MSIToolbox [6] improves this by allowing the measurement of the average m/z intensity values for selected ROIs, but does not provide a way for interactive comparisons between regions. Schramm et al. [7] introduced *imzml*, an open-source format for MSI. This is useful for decoupling the computational methods applied to MS images from the generating hardware. Finally, Metaspace [8] is a recent online MSI repository that allows users to upload MSI data with metabolite annotations. It has support for basic MSI visualizations as grayscale images with the sum of the total intensities for each pixel. However, none of these software solutions enable quick user interaction with the images, beyond clicking on image points. For example, if two interesting regions are observed in a visualization, investigating the difference between them would require saving regions of interest, writing custom analysis code for statistically comparing the ROIs, and then performing a visual verification of the mapped m/z values. In practice, this is rarely done.

We therefore identify two main limitations with current MSI software that we focus on in this study:

- No existing MSI software allows interaction with MSI data and mapping m/z values for different image regions, beyond a trivial selection of image pixels.
- There is no comprehensive support for advanced, modern MSI visualization methods.

1.3 Previous work for visualizing MSI data

Previous work for visualizing high-dimensional MSI data mainly relied on different clustering methods such as K-means [9, 10] or dimensionality reduction (DR) methods such as UMAP [11–14]. In clustering methods, each image pixel is assigned to a cluster and is colored by the assigned cluster color. In DR methods, individual features of the transformed data with reduced dimensions can be visualized as images, resulting in multiple summary images of the data. If the reduction is performed in three dimensions (3D), the dimensions can be interpreted as Red/Green/Blue (RGB) channels and visualized as a color image.

UMAP [15], a modern and widely used DR method, has been introduced for MSI in [16]. This study highlights the effects different distance metrics have on the resulting visualization. Additionally, the use of a *Self-Organized Map* [17] has been proposed in [18], along with the introduction of the distance correlation metric for evaluating MSI visualizations. The need to evaluate MSI data using different distance functions beyond Euclidean distance was also emphasized.

The distance correlation metric introduced in [18] is an application of a global structure preservation metric for dimensionality reduction. Global structure preservation is concerned with maintaining relationships between groups of points in the high-dimensional data and the transformed low-dimensional data. For example, if we have three points, where two are close to each other and the third is far from both, the global structure is preserved if the relationship is retained in the low-dimensional space.

Global structure preservation is especially essential for MSI. Consider the case of small anatomical structures such as healthy neurons or pathological β-amyloid (Aβ) plaques in the brain cortex. These structures have a distinct m/z content compared to other cortical regions but may also share some m/z features. For example, Aβ-plaques in the brains of Alzheimer's disease (AD) mice accumulate m/z values from the surrounding cortex layers. To ensure these small structures are visible in MSI visualizations, we need to preserve the global structure of the high-dimensional data, maintaining contrast between structures and the surrounding areas even when shared m/z content exists. Detail-rich anatomical regions, such as the brain stem, present another challenge. Many m/z features are shared across different sub-regions and nuclei, which can bias visualization methods to assign them similar colors. In extreme but common cases, sub-regions become indistinguishable for the viewer. This lack of differentiation makes it difficult for MSI users to annotate and explore sub-regions. Strong global structure preservation is essential for revealing these internal variations and ensuring that sub-regions do not appear overly homogeneous.

The choice of distance metrics is another critical factor [16, 18]. Some anatomical or functional differences manifest through numerous small changes in m/z values, such as in the brain's white matter or cortical sub-locations (layers). Others are defined by changes in only a few specific m/z values, as seen with Aβ plaques or tumor clones. Different distance metrics may be needed to highlight specific structures effectively.

Many dimensionality reduction methods could be applied to MSI data. Classical methods include PCA [19], FastICA [20], Isomap [21], and NMF [22]. Modern methods include UMAP [15], t-SNE [23], and techniques designed to enhance global structure

preservation, such as Trimap [24], PaCMAP [25], and Phate [26]. PCA [19] is a gold-standard method for global structure preservation. It finds a linear projection that maximizes data variance. Distances between points on the principal component hyperplanes are fully preserved, ensuring global structure retention. However, when reducing MSI data to just three components from thousands of features, PCA may distort the global structure for points far from the principal components. PaCMAP [25] is a recent method that aims to improve global structure preservation in UMAP. Like UMAP, it creates a graph structure with edges between different points, but it expands the edges to include pairs of points that are far apart. This approach indirectly enhances global structure preservation.

Despite these advancements, no previous work has been done to design MSI-optimized visualization methods. Such methods would need to address MSI-specific challenges, including the robustness of different distance metrics, the extremely high dimensionality of MSI (often with thousands of features per pixel), and the necessity for sufficient computational resources. Additionally, visualizations must be evaluated based on their effectiveness in revealing specific or even hidden details within MS images.

1.4 *MSI-VISUAL*: An interactive framework for rapid exploration and mapping of MSI data

To address these gaps in MSI, we propose *MSI-VISUAL*, an open-source framework for extracting, visualizing, and interacting with MSI data. A flow chart describing *MSI-VISUAL* is given in Figure 1.

MSI-VISUAL simplifies MSI analysis by providing an easy-to-use interface for selecting and comparing regions of interest, mapping significant m/z values, and generating multiple visualizations, reducing the time and effort required for MSI interpretation.

First, users create a data pipeline to extract the raw MSI data (either in the open source imzml-format [7] or with additional current support for the Bruker hardware format) to a fixed shape tensor of width $rows \cdot cols \cdot N$ in the open source .npy format. N, the number of features per pixel, is determined by the total range of m/zs in the scan and the number of bins allocated for each m/z unit of 1 Da. Multiple visualizations are created out of the high-dimensional tensor to get a multi-visualization panel, where different visualizations may be better for revealing details in different parts of the image.

Then, users interact with the visualizations by selecting ROIs. *MSI-VISUAL* supports several modes for ROI selection that allow rapid interaction with the MSI data:

1. **Manual annotation of regions.** In this mode, the user annotates polygons around ROIs in different visualizations of the same MSI scan.
2. **Automatic segmentation and clicking on selected regions.** In this mode, first, the visualizations are spatially segmented. We use the classic Felzenszwalb-Huttenlocher spatial segmentation method [27] and convert the MSI visualization to sub-regions. These regions are then selected by clicking on them.

3. **Point selection.** In that mode users click on points, to focus on smaller structures.

Once ROIs are selected, *MSI-VISUAL* compares their m/z content and searches for m/z values that are sufficiently different between the ROIs. These m/z values are then displayed as *SpecQR*, a new graphical visualization that aggregates the m/z spatial images and displays noticeable m/z values on a 2D grid ('QR code'). The aggregation serves as a visual verification for the comparison - if the two regions are separable, there will be a strong visual difference.

Then the users may click on individual m/z values to obtain visualizations of single m/z images (often termed ion images), for a visual verification that these m/z values correspond to the pattern that is searched for.

Fig. 1 The *MSI-VISUAL* **workflow** consists of different stages: MSI data extraction and visualization, interactive exploration with visualization and m/z mapping.

1.5 Novel visualizations designed for MSI

1.5.1 Robust visualization methods for high global structure preservation in MSI

We propose *Saliency Optimization* (SALO) and *Spearman Correlation Optimization* (Spear) to address the next two MSI visualization requirements.

i) High global structure preservation for revealing anatomical details in sub-regions. This is achieved by direct optimization of the global structure preservation. These methods sample a group of representative reference points in the image and optimize an objective on the ranks of the distances of all image pixels from the reference points in the high-dimensional data and the visualization. Both *Spear* and *SALO* rely on a differentiable approximation of the ranking operation [28]. The *Spear* visualization optimizes the Spearman rank correlation [29]. The *SALO* visualization optimizes a new objective we term *Saliency*, which requires ranks of point pair distances in the visualization (in the LAB color space to improve perceptual differences) to be at least as high as in the MSI data to ensure detail is preserved.

Compared to PCA, our methods are non-linear and non-parametric and are able to model complex relationships in the data compared to a linear projection. Compared to UMAP [15] which focuses on local point neighborhoods and local structure preservation, our methods focus specifically on global structure preservation. The idea in PaCMAP [25] of sampling pairs of points that are far from each other is somewhat related to our sampling of reference points. However, unlike PaCMAP, we ***directly*** target global structure by maximizing the distance correlations, therefore optimizing the global structure preservation considering relationships between multiple points at once. This achieves higher global structure preservation: on the mouse brain dataset, as shown in the results section in Figure 2, *SALO* receives an average correlation of approximately 0.8 and PaCMAP approximately 0.45.

Furthermore, since we avoid the graph structures in UMAP and PaCMAP, our methods are simpler to implement, have a lower memory footprint, and therefore are more scalable. On some larger MS images, we were not able to run UMAP even on very strong hardware.

ii) Robustness to difference distance metrics. The Euclidian (and similarly the Cosine) distances are sensitive to changes in many m/zs at once. The L_∞ distance targets large changes in a single m/z. *SALO* and *Spear* achieves robustness to both by fusing the Euclidean and L_∞ norms by taking the maximum of their ranks.

1.5.2 Lightweight visualization methods for large MSI datasets

Next, our goal was to design visualizations that are extremely fast to compute and can handle large datasets that other methods fail on because of the extensive computational requirements. Furthermore, there is a common need for creating thumbnails or quick initial visualizations in MSI viewers. In Metaspace [8] this is done with total ion current (TIC) following by colorizing the image with a color scheme, however, TIC has no ability to distinguish between image regions.

The motivation for these visualizations is the observation that on our mouse brain dataset, in most image pixels, there is a small number of dominating m/zs with high

intensities (with roughly 75% having one dominating m/z, 20% having two m/zs, and the remaining three, or more). The proposed *TOP3* visualization simply measures the three largest intensities in each pixel (and then after outlier removal, converts these from the LAB color space into RGB). The proposed *Percentile-Ratio* (PR3D) visualization is a generalization of this and measures ratios of fixed pre-determined percentiles of the intensities.

As shown in the results section in Figure 3, both of these methods are computationally lightweight, with low memory footprints and are hundreds of times faster to compute. This enables faster computing of MSI visualizations for scans with larger computational requirements (either more pixels or more m/zs) that other methods fail on.

2 Results

2.1 Evaluating the proposed MSI visualizations

2.1.1 Quantitive evaluation of MSI visualizations

To quantitatively evaluate the performance of our proposed MSI visualizations, we benchmarked them on two datasets representing different scenarios, using well-established metrics already used before for MSI (Figure 2). We benchmarked against modern methods: UMAP [15], PaCMAP [25], Phate [26], t-SNE [23], TriMap [24] as well as classic methods: PCA [19], NMF [22], Isomap [21] and FastICA [20].

The newly proposed *SALO* visualization performs better than all existing methods on all datasets, and also for all individual images (i.e., not only on average). The next sections describe the evaluation metrics, the datasets, and the quantitive results.

Fig. 2 Benchmarking of MSI visualization methods. The results are sorted from higher (better) to lower. The proposed new methods are colored in black. The proposed *SALO* method performs better than all existing methods on all metrics and datasets.

2.1.2 Evaluation metrics

As the evaluation metric, we use the distance Pearson correlation metric, which measures the alignment between pairwise distances in the high-dimensional and low-dimensional spaces. This approach has been previously applied to MSI visualization evaluation [18] and is a well-established metric for evaluating dimensionality reduction [30][31]. Following prior work in MSI visualization [18] and to assess performance under different distance metrics corresponding to different biological scenarios of m/z differences between regions, we computed correlations using both Cosine and L_∞ distances.

In addition, in the extended data, we report benchmarking with similar results for the Spearman rank correlation metric (Figure 19), and our own *Saliency* metric (Figure 18) that have not been used before for evaluating MSI, which may be meaningful for future comparison.

2.1.3 Benchmarking datasets

We evaluated the new visualization methods on two datasets:

i) A mouse model brain dataset consisting of four brain hemisphere scans at 20µm resolution. This dataset represents controlled, high-quality MSI data, where we took special care in sample preparation and the measurement protocol.

ii) A diverse collection of six MSI images from the Metaspace repository, covering various tissue types, low and high resolutions, different numbers of features,

and different m/z ranges. Some scans were created with peak selection, reducing the number of features. Some image files are huge, making them computationally challenging to process.

Only our proposed methods (*SALO*, *Spear*, *TOP3*, and *PR3D*), along with PaCMAP, PCA, and NMF, were able to run on the full-size images using a high-performance machine with 320GB of RAM. To facilitate comparisons with all methods on the second dataset, we subsampled every second pixel from the large kidney (Table 2, dataset 4) and lung (Table 2, dataset 6) MS images. The characteristics of the used datasets are summarized in Tables 1 and 2.

Image Size	Features	Tissue type	Resol.	Size
286×262	5255	Mouse brain hemisphere (hA7ko)	20µm	1.5GB
331×248	5255	Mouse brain hemisphere (APPPS1 × hA7ko)	20µm	1.7GB
314×294	5255	Mouse brain hemisphere (APPPS1)	20µm	1.9GB
314×285	5255	Mouse brain hemisphere (C57BL/6J wild-type)	20µm	1.8GB

Table 1 Dataset characteristics for the mouse brain benchmarking dataset. The mouse models were recently described in [32].

Link	Image Size	Features	Uploader	Tissue type	Resol.	Size
Dataset 1	62×160	4505	Manuel Liebeke (MPI Bremen)	Bacterial vesicles	100µm	175MB
Dataset 2	135×199	1755	Schwaiger-Haber et al. [33]	Brain with tumor (mouse)	50µm	184MB
Dataset 3	104×108	4285	Tingting Lu	Ovary (mouse)	20µm	188MB
Dataset 4	550×824	5255	Brittney Gorman (PNNL)	Kidney (human)	5µm	9.3GB
Dataset 5	521×417	10505	Lars Gruber (HS Mannheim)	Spinal cord (mouse)	5µm	8.9GB
Dataset 6	772×596	16505	Dušan Veličković (PNNL)	Lung (human)	20µm	30GB

Table 2 Dataset characteristics for the diverse benchmarking dataset with slides from Metaspace [8].

2.1.4 Benchmarking results

The averaged results across images are presented in Figure 2. Our proposed methods, *SALO* and *Spear*, outperform all existing approaches on both datasets. Additionally, *TOP3* achieves the best performance among existing methods for the L_∞ distance. *TOP3* and *PR3D* show competitive performance, with *TOP3* consistently outperforming UMAP while maintaining lower memory requirements and significantly faster computation on larger images. However, *PR3D* struggles with the diverse dataset, highlighting the need for future work on efficient hyperparameter auto-tuning of the percentiles used in *PR3D*. Beyond the reported average results, we note that *SALO* achieves the highest performance across all metrics for every individual image in both datasets.

2.1.5 Visualization scalability for processing large images

Fig. 3 Benchmarking MSI visualization memory and speed on large images. The proposed, new methods are shown as black bars. *TOP3* is 600× faster to run and consumes only 2**M**B RAM compared to 25**G**B for NMF and 17**G**B for SALO and PaCMAP. Methods that could not run due to resource needs (>320GB RAM) are marked with ✕.

We evaluated the scalability of our methods for large MSI datasets by testing them on the human kidney biopsy sample (Table 2, dataset 4) in the assembled Metaspace dataset. The size of this image is 550×824 pixels with 5,255 features, resulting in a 9.3GB .npy file. We assessed memory usage and runtime, as shown in Figure 3. Despite using high-performance hardware with 320GB RAM, several methods — including UMAP — fail to complete due to excessive memory requirements, indicated by a ✕ in the figure. *TOP3* demonstrates exceptional efficiency, running over 600× faster than PaCMAP while consuming only 2MB of additional memory, compared to approx. 17.5GB required by PaCMAP and *SALO* (approximately twice the image size). *SALO* measures two types of distances from each image pixel to the reference points (N=1,000 by default), resulting in a memory consumption of Rows x Cols x 1000 * 2 floating point values. A lower number of reference points could be traded off to fit memory.

These results highlight the practical advantages of *TOP3*. While it achieves competitive global structure preservation, it does not match the performance of *SALO*. However, its high speed and minimal memory footprint make it particularly well-suited for fast exploring of large MSI datasets.

2.1.6 The proposed visualizations better reveal histomorphological details

To be able to assess the practical improvements of the new visualizations, we needed to analyze examples that include small known histomorphological details.

i) Detecting circumscribed small details. Aβ-plaques in our mouse models [32] are detectable with immunohistochemistry and appear in different sizes and at different locations. They are very circumscribed, highly pathological structures and accumulate numerous lipids [34]. Figure 4 demonstrates that detecting these histopathological changes is better facilitated using the new visualizations. The data used to assess the detection is a lipidomics MSI scan of a mouse brain hemisphere (bregma –2.06 mm) of a 100-day-old mouse with numerous small and large Aβ-plaques.

Whereas UMAP detects only very extensive lipid accumulations in the center of large Aβ-plaques (dense cores), the proposed visualizations *SALO*, *Spear*, *TOP3*, and *PR3D* highlight numerous contrasted, also small Aβ deposits (green arrows for detected plaques, red arrows for non-detected plaques).

Fig. 4 Improved detection of small pathological changes in mouse brain. The detection of histopathological, disease-related objects is better facilitated using the new visualizations (*TOP3*, *PR3D*, *SALO*, and *Spear*). Most Aβ-plaques are missed in UMAP (marked with red arrow).

Parenchymal organs, such as the kidney, have a defined histomorphological appearance with several substructures and sub-regions. The kidney consists of many functional units each consisting of a glomerulus in the cortex and the connected tubulus system in cortex and medulla. Figure 5 shows two visualizations highlighting these details. *SALO* clearly shows the different parts of the tubulus system (proximal and distal tubuli, collection tube) as well as the numerous inflammatory cells between the tubuli and around the glomeruli. The number of glomeruli is reduced and those few that are remaining are strongly pathologically changed. *SALO* also enables clear detection of remaining glomerulus capillaries and inflammatory lymphocytes.

Fig. 5 Human kidney biopsy example with an extensive glomerulonephritis (inflammation). *SALO* visualization is generally more clear in revealing the organ structure than previous methods. a-d) show two examples of glomeruli for the PaCMAP and *SALO* visualizations. The arrows indicate examples that are better detectable in *SALO*: 1) lymphocytes in a glomerulus capillary, 2) proximal tubulus vs. distal tubulus, 3) lymphocytes in a tubulus.

ii) Detecting small structural details in pre-defined anatomical regions.

For research on disease pathomechanisms and for treatment development, the detection of subtle anatomical or functional details within known anatomical regions is of exceptional importance. Figure 6 demonstrates an example of the advantages of the proposed visualizations to highlight such subtle details, e.g. sub-regions or functional domains of the Substantia nigra (SN). Several SN sub-regions representing functional subdomains can be detected visually using the new visualizations. Using UMAP for DR these sub-regions aren't detectable at all. Here, we only detect a homogeneous dark blue area whereas the new visualizations enable the clear detection of sub-regions.

Fig. 6 Improved detection of Substantia nigra (SN) sub-regions in mouse brain. SN sub-regions can be clearly detected using the new visualizations. UMAP only shows a homogenous region. The green arrows pinpoint to examples of sub-regions detectable in *SALO*, *Spear*, *PR3D*, and *TOP3*.

Another example of much-increased detail richness is seen in Figure 7. It demonstrates that sub-nuclei of the thalamus region can be clearly visualized using the newly proposed visualizations. Small thalamic nuclei are especially pronounced using *SALO* and *TOP3* visualizations. *SALO* was used to annotate examples of thalamic sub-nuclei (dotted lines, arrows). The annotation mask was then overlayed with the other visualizations to enable comparisons and orientation within the thalamus region. UMAP shows much reduced details within the thalamus annotation mask. *PR3D* and *TOP3* reveal some but not all thalamic nuclei.

Fig. 7 Improved detection of thalamus nuclei. Our proposed *SALO* and *TOP3* reveal more sub-region details in the annotations compared to UMAP. UMAP shows a homogenic area inside the thalamus region. The red arrows point to thalamus nuclei that are not distinguishible in UMAP, but in *SALO* we can identify and annotate them.

iii) Detecting functional details. We also wanted to assess the ability of the proposed visualizations to reveal hidden functional details in mouse brains, e.g. hippocampal interconnections between neuronal groups. In contrast to the circumscribed, pathological Aβ-plaques, these structures are physiological, larger, and may not appear as pronounced. Extended Data Figures 16 and 17 show an overview of UMAP, *SALO*, *Spear*, *PR3D*, and *TOP3* visualizations revealing different levels of anatomical details of the hippocampal interconnections related to global data structure preservation.

iv) Detecting single neurons. The detection of single cells, e.g. large neurons of the spinal cord's grey matter, is further explained in the interactive mapping section (2.2.5).

2.2 Interactive m/z mapping on biological examples using *MSI-VISUAL*

Determining the spatial distribution of m/z values of metabolites, lipids, and peptides specific for model organisms, anatomical regions, experimental outcomes, etc. is an important task in research. Furthermore, disease-related marker m/zs could be used for diagnostics but also for treatment stratification, e.g. MSI profiles of malignant tumors.

To test the general use of *MSI-VISUAL* to map m/zs on different biological samples with different m/z ranges and spatial resolutions, we analyzed several different cases:

1. Mapping of lipids of anatomical brain regions in mice.
2. Mapping Metaspace repository examples: brain tumor and kidney.
3. Mapping human diagnostic brain samples with neuropathological features.
4. Mapping mouse brain stem nuclei and white matter tracts.
5. Mapping single cells.

2.2.1 Interactive mapping of mouse brain shows region-specific lipid signatures

Here, we demonstrate the usage of the entire *MSI-VISUAL* pipeline for fast interactive mapping of m/z values of different brain regions of a lipid MSI scan of an AD mouse model with pronounced β-amyloidosis [32]. *SpecQR* examples of the m/z distribution are given in Figure 8. We selected and mapped the cortex, hippocampus, white matter, and brain stem. Each of these regions has a unique lipid signature. The top 5 m/z values are displayed for every region. We show individual m/z images of detected m/zs for verification. Previous work [35] used registration of MSI to MRI to be able to map m/z values belonging to these brain regions. Using *MSI-VISUAL* we are similarly able to map the brain regions, however without the need for registration to external data modalities such as MRI. We achieve this almost instantly in a fraction of the time compared to the process proposed by [35], by simply selecting regions in the software.

We also map Aβ-plaques simply by selecting them in the MSI viewer. Of notice is that some m/z values, such as 885.6 Da, highlight the plaques in addition to cortex regions, showing that the Aβ-plaques absorb and concentrate lipids of the cortex neurons. Other m/z values like 1179.8 Da are more unique to the Aβ-plaques in the cortex and are only found outside the cortex, e.g. hippocampus. This is the first time, m/z values for plaques are determined from within MSI images, in contrast to searching for known, pre-determined, plaque m/z values.

Fig. 8 Using *MSI-VISUAL* for interactive mapping of the m/z components of a brain lipid MSI scan of an AD mouse model with Aβ-plaques. **Upper image**: SpecQR summary images for a) cortex (CTX), b) hippocampus (HIP), c) brain stem (BST), d) white matter (WM), e) lower, and f) upper cortex regions (CTX). **Bottom image**: g) Aβ-plaques have a distinct lipid m/z signature with 599.4, 885.6, 886.6, 1179.8, and 1180.8 being the 5 highest m/z values in this example. The SpecQR summary image and the images of the individual top m/zs are used for verification (in brown). h) An image with numerous cortex (CTX) plaques (arrow 1) containing the plaque-specific m/z 1179.8 Da (ganglioside GM3 [34]) which is also found in hippocampal connections (arrow 2).

2.2.2 Interactive mapping of metabolomics and lipidomics datasets from Metaspace

To test the general applicability of *MSI-VISUAL*, we wanted to assess MSI data generated by others available on Metaspace [8]. The results for m/z mapping with a human kidney biopsy (lipidomcis, 300-1350 Da) and a mouse brain with implanted brain tumor (metabolomics, 50-400 Da) are shown in Figure 9. These examples show the calculation of regions-specific m/z values. The kidney biopsy of a glomerulonephritis patient nicely highlights the anatomical structures (glomeruli and tubuli) but also the extensive roads of inflammatory cells and their unique m/z pattern. The brain tumor example highlights the tumor-specific metabolite m/zs.

Fig. 9 Demonstrating *MSI-VISUAL* on images from the Metaspace repository. Upper panel: Mapping lipidomics m/zs 300-1350 Da of different regions in a human kidney biopsy (Table 2, dataset 4), aggregated images of the 5 most specific m/z for a) kidney tubuli (green arrows) and b) inflammatory cells (orange arrows) around glomeruli and between the tubuli. c) 2-channel image of tubuli (red) and inflammation (green). Bottom panel: d) Mapping metabolomics m/zs 50-400 Da of a transplanted brain tumor (yellow arrow) into a mouse brain from [33]. The *SpecQR* shows the specific metabolite pattern of the lesion. Below, ion images of the 5 best m/z and a differential red-green image (tumor vs. all) are shown.

2.2.3 Interactive mapping of human brain samples to diagnose neuropathological lesions

Post-mortem human brain tissue of a patient with vessel-related dementia and AD was investigated using peptide MSI. The data have been pre-processed and are peak-selected for the 500 best peaks [36]. We wanted to automatically identify diagnostically relevant neuropathological lesions in these brain sections. The example in Figure 10 shows the identification of Aβ-deposits related to AD and additionally, the severe capillary leakage related to vascular dementia (VD). We selected the segmented ROIs according to their color (green or red/blue), and MSI-visual generated the output of the 5 most significant m/z values for each ROI. The identification of these specific m/z values enabled the correct identification of the ongoing neuropathological processes and their tissue locations in this mixed dementia case.

Fig. 10 Visualization and identification of intracerebral lesions to diagnose neuropathological conditions by assessing their specific m/z values. a & b) NMF visualization (2000 iterations) presented with two different colorizations, RGB and BGR. The box and red/blue arrows mark examples of numerous smaller, partially accumulating/aggregating structures, the green arrows highlight larger confluent structures. These structures were selected for further m/z characterization and identification of their pathological origin. c) the m/z fingerprints and SpecQRs reveal extensive capillary leakage and d) larger Aβ depositions with Aβ1–40 (4332.2 Da), Aβ1–38 (4129.8 Da). e & f) presentations of the three most significant m/z values (1–3) and the aggregated mean images (aggr) – e) highlights severe capillary leakage as seen in vascular dementia, f) shows β-amyloid pathology as seen in Alzheimer's disease, thus the diagnosis here is a mixed-type dementia. Pride repository dataset PXD049325, peak-selected, peptide MSI, case S2-SAD [36].

2.2.4 Mapping mouse brain nuclei and white matter tracts

As an additional example, *MSI-VISUAL* facilitates the identification and characterization of functional groups of neurons (nuclei) and their connections (white matter tracts). The interactive visualizations enable the mapping of specific m/z values to selected brain structures, providing high functional resolution, as demonstrated for hippocampal neurons and connections (Figure 11).

Fig. 11 Visualization and functional segmentation to detect specific m/z values for anatomical and functional sub-regions. a) *TOP3* segmentation of brain hemisphere lipidomics MSI data revealed detailed structures, including ROI1 (substantia nigra, SN) and ROI2 (cerebral peduncle, cpd), which were compared using *MSI-VISUAL*. b) ROI1 (SN) was identified as a neuronal region, with m/z values corresponding to other high-neuron areas like the hippocampus (HIP). c) ROI2 (cpd) was correctly identified as a white matter tract, with no overlap with SN. d) Overlay of the top five m/z values from ROI1 (red) and ROI2 (green), with *SpecQR* diagram showing m/z distribution from 300–1350 Da, confirming distinct m/z profiles between grey and white matter.

We applied *MSI-VISUAL* for functional segmentation to map m/z values associated with anatomical and functional sub-regions in brain hemisphere lipidomics MSI data. The *TOP3* method revealed structures such as the substantia nigra (SN) and cerebral peduncle (cpd), which were used as regions of interest (ROIs) for m/z comparison.

ROI1 (substantia nigra) was characterized by neuronal content and connections, while ROI2 (cerebral peduncle) represented white matter. The m/z values in ROI1 also corresponded to areas rich in neuronal content, such as the hippocampus and its connections. ROI2 was correctly identified as white matter, with no overlap in m/z values between the grey and white matter regions. *SpecQR* diagrams confirmed these distinct m/z distributions.

2.2.5 Mapping single cells

We also wanted to assess the possibility of detecting and analyzing single cells. We chose spinal cord neurons as examples and used an MSI slide with 5μm resolution (Table 2, dataset 5) as shown in Figure 12. Selecting the specific ROI of a single neuron was performed 3-times. Each time the *SpecQR* showed signal enrichment in the neurons and partially in the surrounding grey matter of the spinal cord.

Fig. 12 Example of single-cell detection of spinal cord neurons in a mouse. a-c) show three independent rounds of ROI section (single neuron) and the calculated specific m/z values and SpecQR for these spinal neurons. d) shows a red-green overlay image for the neurons (red channel) and the remaining regions (green). e & f) show two ion images of specific m/z examples (883.6 and 790.6 Da). The red dotted line delineates the grey substance of the spinal cord. The green arrows pinpoint to 2 examples of detected neurons. ant - anterior, post - posterior.

2.3 Applications of *MSI-VISUAL* for biological interpretation of mapped m/zs

We present two examples demonstrating how *MSI-VISUAL* facilitates biological insights through the visualization of mapped m/z values. This process is enabled by the entire pipeline: the new visualizations reveal sub-regions, which can then be interactively investigated using the m/z mapping functionality of *MSI-VISUAL*.

Figure 13 illustrates the identification of a lipid that is enriched in a sub-region of the Substantia nigra (SN) in an ABCA7-deficient mouse, highlighting the localized effects of ABCA7 loss. To explore region-specific lipid differences in this ABCA7-knockout mouse model [32], we compared UMAP and *SALO* visualizations for mapping m/z values across brain sub-regions. While *SALO* successfully distinguished two sub-regions within a selected region of interest (ROI1), UMAP failed to resolve these differences. Within ROI1, we identified a peak at 722.6 Da, which was strongly enriched in this region. A second peak, at 885.6 Da, exhibited a weaker, though noticeable, preference for ROI1. The 722.6 Da lipid, detected in negative mode using a NEDC matrix, is a candidate for several lipid species, including phosphatidylethanolamine (PE 34:1), phosphatidylserine (PS 32:0), and ceramide (Cer 42:2;O3 [M+OAc]-).

Figure 14 demonstrates the determination of m/z values associated with Aβ-plaques. To investigate the lipid composition within these plaques, we examined the spatial distribution of m/z 885.6 (PI 38:4) using *MSI-VISUAL*. This lipid, known to be abundant in neuronal connections in the brain [37, 38], was also detected by us in spinal cord neurons. When comparing visualization methods, *TOP3* successfully highlighted numerous Aβ-plaques, while UMAP failed to resolve them.

Further analysis using *SpecQR* diagrams of both the spinal cord and brain revealed that several m/z values were shared between spinal cord neurons and cortical Aβ-plaques. Ion images of m/z 885.6 in two AD mouse models revealed its presence in large spinal neurons, Aβ-plaques, and hippocampal neuronal connections. These findings suggest that m/z 885.6 is not only a physiological component of neuronal structures but also accumulates in Aβ-plaques, likely due to the degradation of neuronal connections within these pathological structures.

Fig. 13 Index lipids mark region-specific neuronal differences in an ABCA7-knockout mouse. (a, b) SALO identified two distinct sub-regions within ROI1 (green arrows), which were not resolved by UMAP (red arrows). A lipid at 722.6 Da was strongly enriched in this region.

.

Fig. 14 *MSI-VISUAL* **reveals shared lipid content in spinal cord neurons and Aβ-plaques**. (a-d) Visualizations of spinal neurons and Aβ-plaques; UMAP does not highlight plaques (c), whereas TOP3 successfully detects them (d). (e-f) *SpecQR* diagrams show m/z values shared between spinal neurons and cortical Aβ-plaques. (g-i) Ion images of m/z 885.6 in the spinal cord and brain of AD mouse models reveal its presence in spinal neurons (g) and hippocampal neuronal connections (h), as well as its accumulation in Aβ-plaques (h-i). Red arrows indicate plaques missed by UMAP.

3 Discussion

Our results demonstrate that both *SALO* and *Spear* consistently outperform other visualization methods across the two benchmarked datasets. *SALO* notably performs better than all existing methods on all images individually. Compared to UMAP [15] previously used for MSI [16], *SALO* improves by a large margin (approx. 0.3 and 0.4 for UMAP on both datasets and above 0.8 for *SALO*). The low performance of UMAP means that UMAP should be used with caution in MSI, potentially measuring its performance when applied. *SALO* also improves upon PaCMAP [25] which was created to improve UMAP's global structure preservation, and upon PCA [19], a classic gold standard method for this. As noted in the introduction, we attribute this to the direct optimization of global structure in *SALO* compared to the indirect optimization used in PaCMAP, and to the complex non-linear transformation we used compared to the linear projection used by PCA. *SALO* and *Spear* also show the utility in combining different distance functions, resulting in robustness to both distance functions, another need in MSI as discussed in our introduction and noted before in [16, 18].

In contrast, *TOP3* is surprisingly competitive despite its simplicity. While *TOP3* performs better than UMAP [15] it does not achieve the same level of performance as *SALO* and *Spear*. Despite this, *TOP3* stands out due to its low computational requirements, positioning it as a practical tool for MSI analysis, particularly in resource-constrained environments, where methods such as *SALO* or UMAP are not able to run. Notably, because *TOP3* operates on pixels independently, a future optimization could enable partial loading of MSI data, allowing for visualization of image blocks without requiring the entire dataset to be loaded at once.

The *PR3D* visualization, however, did not perform well on all Metaspace datasets. This outcome was partially anticipated, as the hyperparameters for *PR3D* were specifically tuned for the mouse brain dataset. Future work will focus on developing methods for the fast, automatic selection of optimal *PR3D* parameters to enhance its performance across different datasets.

Several limitations of our study must be acknowledged. Although our visualizations emphasize high global structure preservation and reveal anatomical details that assist in annotation, local structure preservation was not the primary focus of this study. For instance, in the case of the *TOP3* visualization, unrelated regions are assigned the same color (e.g., red), which can compromise local structure preservation. Addressing local structure preservation in MSI and its use cases is a potential avenue for future research.

Another limitation of our visualization benchmarking is that for all methods we use the default hyperparameters (excluding UMAP where we applied it twice with different distance functions). It is possible that searching for optimal hyper-parameters for the specific benchmarks would improve their performance as well as the performance of our methods, although it is not practical and we expect visualizations to be created with a single set of parameters in most cases. Since this is extremely computationally intensive, we leave this for future studies.

We also anticipate potential improvements for *SALO*. A major limitation is its reliance on a fixed set of reference points. While re-sampling these reference points could enhance the visualizations, it may come at the cost of increased convergence time

and higher computational expense. Thus, future work will explore ways to balance performance improvements with computational efficiency.

Another challenge with the *MSI-VISUAL* pipeline is its dependence on images extracted using fixed bins and the summation of m/z values within those bins. Extracting m/zs with higher precision thus requires using a larger number of bins, and more memory. Although this approach is standard in MSI, future advancements could explore methods to optimize or even bypass this process, allowing mapping m/z values with full precision.

Finally, *MSI-VISUAL* offers a rapid and interactive platform for mapping and exploring MSI data, as illustrated in the biological examples section. By utilizing *MSI-VISUAL*, we were able to identify and select sub-regions within our visualizations and perform efficient m/z mapping, revealing new biological insights into the shared lipid content of neurons, plaques, and substantia nigra sub-regions. These capabilities underscore the potential of *MSI-VISUAL* for applications that demand fast results, such as for clinical diagnostics or in research settings requiring swift exploration of MSI data.

4 Methods

4.1 The lipidomics MSI mouse brain dataset

The mouse brain MSI data files (lipidomics) were generated using frozen hemispheres of four 100-day-old animals in the C57Bl/6J background. Two mice expressed combined human mutated APP and PS1-mutated transgenes (APPtg) for the generation of Aβ plaques [39–41]. Frozen mouse hemispheres were coronary cut in 10µm thick sections using a cryomicrotome (CM1950, Leica, Nussloch, Germany) and placed on a single Intellislide (Bruker Daltonics, Bremen, Germany). The section was dried in a vacuum chamber and sprayed with N-(1-Naphthyl) ethylenediamine dihydrochloride (NEDC) matrix dissolved to 7 mg/ml in 70% methanol using a HTX3+ TM Sprayer (HTX technologies, Carrboro, NC, USA) as described by Andersen et al. [42]. Matrix-assisted laser desorption ionization trapped ion mobility mass spectrometry imaging (MALDI-TIMS-MSI) was performed in negative mode using a timsTOF fleX™ mass spectrometer (Bruker Daltonics, Bremen, Germany). Mass calibration was performed using red phosphorous and ion mobility using ESI-L Low Concentration Tuning Mix (Agilent Technologies, Santa Clara, CA, USA) in ESI mode. For MALDI-TIMS-MSI analysis, a mass range of m/z $300 - 1350$ and a range of the inverse reduced mobility (1/K0) of 0.7–1.8 Vs/cm^2 was used. The laser was operated at a frequency of 10 kHz with a burst of 200 shots per pixel with a spatial resolution of 20 μm. Four mouse hemispheres were assessed in a single overnight run. The mouse brain lipidomics dataset files are available at the ProteomeXchange PRIDE repository with the dataset number PXD056609.

4.2 The human proteomics MSI dataset

The human brain MSI data file have been just recently published in [36] and have been downloaded from the PRIDE repository with the access code: PXD049325.

4.3 The metabolomics and lipidomics Metaspace datasets

The Metaspace dataset characteristics are highlighted in Table 2 in Section 2.1.3.

4.4 Data extraction and pre/post-processing information

For the extraction and binning of the mouse brain dataset, we used the *"extract_bruker_tims.py"* script from MSI-VISUAL. For the extraction and binning of the slides from Metaspace, we used the *"extract_pymzml.py"* from msi visual. We used 5 bins per m/z in all experiments, resulting in a resolution of 0.2 Da. During the extraction, m/z values are rounded to their closest bin and are summed there.

When creating the visualizations, we use TIC [43] normalization, which in every pixel divides each m/z intensity by sums of the intensities for that pixel.

In all visualizations, we use outlier clipping and clip the 0.01 and 99.99 percentiles of the data, also known as Winsorizing [44].

4.5 Visualization methods

4.5.1 Motivation for the saliency metric used in *SALO*

We first outline the motivation behind the saliency metric used in *SALO*. This metric serves as an objective function for optimization but can also be used to evaluate MSI visualizations.

An effective saliency metric should satisfy the following key properties:

- **Preservation of important details.** Pairs of points that exhibit large distances in the input space but are mapped to small distances in the visualization should be penalized. This ensures that meaningful structural differences are preserved.
- **Higher weight for rare but significant differences.** Pairs of points with large distances in the input space should have a greater influence on the saliency score. While such pairs may constitute a small fraction of the total, they are critical for saliency. Without this weighting, structures comprising a small percentage (e.g., $< 1\%$) of all pairs might be completely overlooked, yet the overall metric could still appear high. Additionally, points with very small input-space distances can artificially inflate the saliency score, despite not contributing to meaningful structure identification.
- **Integration of Cosine and L_∞ distances.** The metric should accommodate multiple distance functions to capture diverse biological structures. Some structures manifest themselves through a few distinct m/z values with large intensity changes (better captured by the distance L_∞, while others involve coordinated changes across many m/z values with smaller individual intensity variations (better captured by the distance of Cosine). Both types of patterns should be detected.
- **Incorporation of perceptual color differences.** The metric should account for human perceptual biases in color interpretation. For instance, subtle variations in shades of pink may be difficult to distinguish, even if the visualization encodes meaningful numerical differences. A saliency metric that aligns with human visual perception ensures that detected structures are recognizable to users.

4.5.2 The Saliency objective definition

For each pair i, we compute its ranking among the other Cosine and L_∞ distances and take the larger ranking:

$$R_{\max}(X_i) = \max(R_{\text{cosine},i}, R_{L_\infty,i})$$

This ensures that property 3 is satisfied. Moving from raw distance values to ranks allows us to compare values from different distance functions and directly relate input distances to output distances.

For the output, we measure the L_2 distance in the visualization space, specifically in the LAB color space, to account for perceptual differences (satisfying property 4). We then compute the rank $R(Y_{\text{LAB},i})$ of each pair in this color space.

We define Δ-saliency for each pair of points i as the ability of the visualization to rank the output at least as high as the input, with a tolerance of Δ (satisfying property 1):

$$\Delta\text{-saliency}_i = \sigma_\Delta(i) = \mathbb{I}\left(R(Y_{\text{LAB},i}) - R_{\max}(X_i) + \Delta > 0\right) = \begin{cases} 1, & \text{if } R(Y_{\text{LAB},i}) - R_{\max}(X_i) + \Delta > 0 \\ 0, & \text{otherwise} \end{cases}$$

The overall Δ-saliency is computed as the weighted average of $\sigma_\Delta(i)$ across all point pairs, where the weights are determined by their ranks (ensuring property 2):

$$\Delta\text{-saliency} = \sum_i \frac{\sigma_\Delta(i) \cdot R_{\max}(X_i)^\gamma}{\sum_i R_{\max}(X_i)^\gamma}$$

Finally, we define the **Average Saliency** as the mean of the Δ-saliency values for $\Delta = 10\%$, 20%, and 30%.

4.5.3 Saliency Optimization

We approximate the *Saliency* metric, and then optimize for a visualization that maximizes it with stochastic gradient descent (SGD).

We approximate Δ-saliency by sampling fixed reference points to reduce the number of pairs from being quadratic with the number of pixels. *SALO* optimizes over the set P all pairs of each image point, and all the reference points:

$$P = \{(a, b) \mid a \in \text{Image}, b \in \text{Reference points}\}$$

The average *Saliency* metric defined above as the sum over all image pairs, is then instead restricted to the set of pairs P.

The ranking is not differentiable. Therefore we use *soft ranking*, a differentiable approximation of ranking from [28]. We use the default parameters given from the official PyTorch [45] implementation.

Finally, plugging in the soft rank apprximation $\hat{R_{\max}}(X_i)$ to the saliency metric, summing over the set P, and negating for minimization, the saliency loss function which we minimize with gradient descent is:

$$L_{\Delta\text{-saliency}} = - \sum_{i \in \{(a,b) \mid a \in \text{Image}, , b \in \text{Reference points}\}} \frac{\mathbb{I}\left(R(Y_{\text{LAB},i}) - \hat{R_{\max}}(X_i) + \Delta > 0\right) \cdot \hat{R_{\max}}(X_i)^\gamma}{\sum_i \hat{R_{\max}}(X_i)^\gamma}$$

In all experiments, we set $\Delta = 0$ and $\gamma = 2$ in the loss function, based on the hypothesis that this would enhance the detection of rare objects, such as plaques or neurons. However, these values could be further refined. As previously noted, this parameter assigns greater weight to larger distance ranks, which may help in identifying rare objects. A lower value of γ, such as 0 or 0.5, could potentially improve overall performance while reducing effectiveness in rare sub-groups. Further investigations into this trade-off would be valuable in future work.

As a strategy for selecting reference points for optimization, we consider the next options:

- **Lightweight core sets method** [9]. This method is a fast sampling method that samples points for both diversity and coverage of main clusters.
- **Kmeans++** [10], which selects points in a way that maximizes their spread: the first point is chosen randomly, and subsequent points are sampled with probability proportional to their squared distance from the nearest existing centroid.
- **Random sampling**.

Figure 15 shows the effect of different sampling strategies. For a large number of points, the difference is smaller, however, run time will increase (the complexity is proportional to the number of reference points). Cosine distance correlation performs better with Kmeans++, however, this is somewhat expected because the distance function in Kmeans++ is Euclidean.

By default, we use the lightweight core sets method [9] for sampling reference points because of its speed.

For optimization we use the Adam optimizer [46] with a learning rate of 1. The visualization is initialized with a random normal distribution in the range [-5, 5].

Fig. 15 The effect methods for selecting _SALO_ reference points has on visualization quality.

4.5.4 Spearman Rank Correlation Optimization

The Spearman visualization method is similar to $SALO$, however, the objective is to maximize the Spearman rank correlation between the maximum rank of the cosine and L_∞ distances in the high- and low-dimensional data.

This provides a simpler and more natural alternative to the saliency objective, which we use as a baseline to verify its positive impact.

4.5.5 Lightweight visualizations

TOP3 selects the three highest intensity values per pixel and then maps them to the LAB color space. We implement this method efficiently using the Numba Python package.

Our **_Percentile Ratio (PR)_** visualization is a generalization of $TOP3$ that measures how intensities are distributed. We define the PR function, which operates on pixel coordinates i and j for two fixed percentiles p_1 and p_2:

$$\lambda(X_{ij,p_1,p_2}) = \frac{\text{Percentile}(X_{ij}, p_1)}{\text{Percentile}(X_{ij}, p_2)}$$

The L, A, and B channels in the LAB color space are then computed as follows:

$$L_{ij} = \lambda(X_i j, p1, p2) = \lambda(X_{ij}, 99.99\%, 99.9\%)$$
$$A_{ij} = \lambda(X_i j, p2, p3)\lambda(X_{ij}, 99.9\%, 99\%)$$
$$B_{ij} = \lambda(X_i j, p4, p5)\lambda(X_{ij}, 98\%, 85\%)$$

The percentiles p_1, p_2 and p_3 were selected by a grid search on the first image of the mouse brain dataset (containing 5,255 features), evaluating 1,000 predefined combinations within the range $[85\%, 100\%]$. The values of p_4 and p_5 were manually adjusted to achieve the desired pink shade. Future work may focus on making these values differentiable for direct optimization or developing a more efficient search method.

4.6 Interactive m/z mapping methods

4.6.1 Region selection

MSI-VISUAL facilitates the automatic selection of regions by interacting with the visualization. One method for this is based on the approach by Felzenszwalb et al. [27], a classical, efficient, and fast graph-based computer vision algorithm that performs spatial segmentation of images into sub-regions, accounting for both spatial arrangement and visual appearance.

An alternative method supported by *MSI-VISUAL* involves segmentation of color channels without using spatial information. In this case, each of the three color channels (Red, Green, and Blue) in the visualization is segmented into N sub-regions (with $N = 5$ by default, though this can be adjusted). Segmentation is achieved by digitizing each channel into N equally spaced bins in the range $[0, 1]$:

$$\text{Segment}_{ij} = \text{Digitized}(\text{Red})_{ij} \cdot N^2 + \text{Digitized}(\text{Green})_{ij} \cdot N + \text{Digitized}(\text{Blue})_{ij}$$

Here, i and j denote the pixel coordinates, Digitized(image) represents the quantization of an image or region into N bins, and Segment_{ij} refers to the sub-region to which pixel (i, j) belongs.

4.6.2 Region vs. region comparisons

For a pair of selected regions, the Area Under the Curve (AUC) is computed between the two groups. We then perform statistical analysis on the m/z values, selecting those with a p-value<0.05 to identify m/z values that significantly differ between the compared regions.

4.6.3 Region vs. all comparisons

In this mode, the user selects a single region for comparison with the entire image. A complementary mask is created for the selected region, and to speed up the comparison, every second pixel of the complementary mask is sub-sampled.

4.6.4 Point-to-Point comparisons

In this mode, users select pairs of points by clicking on them within the visualization. For each m/z value in the spectra of a point, we compute its rank within the image. The rank difference between the two selected points is then calculated, and the m/z values are sorted by this rank difference. A positive rank difference indicates that a particular m/z value is more prominent in the first point, while a negative rank difference implies that the m/z value is more prominent in the second point.

4.7 Data availability

Our mouse brain dataset files are available at the ProteomeXchange PRIDE repository with the dataset number PXD056609.

4.8 Materials availability

Not applicable.

4.9 Code availability

We release the code as an open-source package at Github: MSI-VISUAL.

References

[1] Shimma S. Mass Spectrometry Imaging [Journal Article]. Mass Spectrom (Tokyo). 2022;11(1):A0102. https://doi.org/10.5702/massspectrometry.A0102.

[2] Alexandrov T. Spatial Metabolomics and Imaging Mass Spectrometry in the Age of Artificial Intelligence [Journal Article]. Annu Rev Biomed Data Sci. 2020;3:61–87. https://doi.org/10.1146/annurev-biodatasci-011420-031537.

[3] Nunes JB, Ijsselsteijn ME, Abdelaal T, Ursem R, van der Ploeg M, Giera M, et al. Integration of mass cytometry and mass spectrometry imaging for spatially resolved single-cell metabolic profiling [Journal Article]. Nat Methods. 2024;21(10):1796–1800. https://doi.org/10.1038/s41592-024-02392-6.

[4] Zhang H, Lu KH, Ebbini M, Huang P, Lu H, Li L. Mass spectrometry imaging for spatially resolved multi-omics molecular mapping [Journal Article]. Npj Imaging. 2024;2(1):20. https://doi.org/10.1038/s44303-024-00025-3.

[5] Bemis KA, Foll MC, Guo D, Lakkimsetty SS, Vitek O. Cardinal v.3: a versatile open-source software for mass spectrometry imaging analysis [Journal Article]. Nat Methods. 2023;20(12):1883–1886. https://doi.org/10.1038/s41592-023-02070-z.

[6] Zhang Y, Wang X, Li M, Guo T, Zhao Z, Zhang X, et al. An easy-to-use graphical user interface for mass spectrometry imaging analysis [Journal Article]. International Journal of Mass Spectrometry. 2023;492. https://doi.org/10.1016/j.ijms.2023.117105.

[7] Schramm T, Hester Z, Klinkert I, Both P, Heeren RMA, Brunelle A, et al. imzML – A Common Data Format for the Flexible Exchange and Processing of Mass Spectrometry Imaging Data. Journal of Proteomics. 2012;75(16):5106–5110. https://doi.org/10.1016/j.jprot.2012.07.026.

[8] Alexandrov T, Ovchinnikova K, Palmer A, Kovalev V, Tarasov A, Stuart L, et al. METASPACE: A community-populated knowledge base of spatial metabolomes in health and disease. bioRxiv. 2019;https://doi.org/10.1101/539478.

[9] Bachem O, Lucic M, Krause A. Scalable k-Means Clustering via Lightweight Coresets [Journal Article]. arXiv. 2017;https://doi.org/10.48550/arxiv.1702.08248.

[10] Arthur D, Vassilvitskii S. k-means++: the advantages of careful seeding. In: SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics; 2007. p. 1027–1035. Available from: https://theory.stanford.edu/~sergei/papers/kMeansPP-soda.pdf.

[11] Hu H, Yin R, Brown HM, Laskin J. Spatial Segmentation of Mass Spectrometry Imaging Data by Combining Multivariate Clustering and Univariate Thresholding [Journal Article]. Anal Chem. 2021;93(7):3477–3485. https://doi.org/10.1021/acs.analchem.0c04798.

[12] Zhang W, Claesen M, Moerman T, Groseclose MR, Waelkens E, De Moor B, et al. Spatially aware clustering of ion images in mass spectrometry imaging data using deep learning [Journal Article]. Anal Bioanal Chem. 2021;413(10):2803–2819. https://doi.org/10.1007/s00216-021-03179-w.

[13] Prasad M, Postma G, Franceschi P, Buydens LMC, Jansen JJ. Evaluation and comparison of unsupervised methods for the extraction of spatial patterns from mass spectrometry imaging data (MSI) [Journal Article]. Sci Rep. 2022;12(1):15687. https://doi.org/10.1038/s41598-022-19365-4.

[14] Buchberger AR, DeLaney K, Johnson J, Li L. Mass Spectrometry Imaging: A Review of Emerging Advancements and Future Insights [Journal Article]. Anal Chem. 2018;90(1):240–265. https://doi.org/10.1021/acs.analchem.7b04733.

[15] McInnes L, Healy J, Melville J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv preprint arXiv:180203426. 2018;https://doi.org/10.48550/arXiv.1802.03426.

[16] Smets T, Verbeeck N, Claesen M, Asperger A, Griffioen G, Tousseyn T, et al. Evaluation of Distance Metrics and Spatial Autocorrelation in Uniform Manifold Approximation and Projection Applied to Mass Spectrometry Imaging Data. Analytical Chemistry. 2019;91(9):5706–5714. https://doi.org/10.1021/acs.analchem.8b05827.

[17] Kohonen T. Self-Organized Formation of Topologically Correct Feature Maps. Biological Cybernetics. 1982;43(1):59–69. https://doi.org/10.1007/BF00337288.

[18] Gardner W, Maliki R, Cutts SM, Muir BW, Ballabio D, Winkler DA, et al. Self-Organizing Map and Relational Perspective Mapping for the Accurate Visualization of High-Dimensional Hyperspectral Data [Journal Article]. Anal Chem. 2020;92(15):10450–10459. https://doi.org/10.1021/acs.analchem.0c00986.

[19] Pearson K. On Lines and Planes of Closest Fit to Systems of Points in Space. Philosophical Magazine. 1901;2(11):559–572. https://doi.org/10.1080/14786440109462720.

[20] Hyvarinen A. Fast and robust fixed-point algorithms for independent component analysis [Journal Article]. IEEE Trans Neural Netw. 1999;10(3):626–34. https://doi.org/10.1109/72.761722.

[21] Tenenbaum JB, de Silva V, Langford JC. A global geometric framework for nonlinear dimensionality reduction [Journal Article]. Science. 2000;290(5500):2319–23. https://doi.org/10.1126/science.290.5500.2319.

[22] Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. Nature. 1999;401:788–791. https://doi.org/10.1038/44565.

[23] van der Maaten L, Hinton G. Visualizing Data using t-SNE. Journal of Machine Learning Research. 2008;9:2579–2605.

[24] Amid E, Warmuth MK. TriMap: Large-scale Dimensionality Reduction Using Triplets. ArXiv. 2019;abs/1910.00204. https://doi.org/10.48550/arXiv.1910.00204.

[25] Wang Y, Huang H, Rudin C, Shaposhnik Y. Understanding how dimension reduction tools work: an empirical approach to deciphering t-SNE, UMAP, TriMap, and PaCMAP for data visualization. J Mach Learn Res. 2021;22(1). https://doi.org/10.5555/3546258.3546459.

[26] Moon KR, van Dijk D, Wang Z, Chen W, Hirn MJ, Coifman RR, et al. PHATE: a dimensionality reduction method for visualizing trajectory structures in high-dimensional biological data [Journal Article]. BioRxiv. 2017;120378. https://doi.org/10.1038/s41587-019-0336-3.

[27] Felzenszwalb PF, Huttenlocher DP. Efficient Graph-Based Image Segmentation. International Journal of Computer Vision. 2004;59(2):167–181. https://doi.org/10.1023/B:VISI.0000022288.19776.77.

[28] Blondel M, Teboul O, Berthet Q, Djolonga J. Fast Differentiable Sorting and Ranking [Journal Article]. arXiv. 2020;https://doi.org/10.48550/arXiv.2002.08871.

[29] Spearman C. The Proof and Measurement of Association Between Two Things. The American Journal of Psychology. 1904;15(1):72–101. https://doi.org/10.2307/1412159.

[30] Lee JA, Verleysen M. Quality assessment of dimensionality reduction: Rank-based criteria. Neurocomputing. 2009;72(7-9):1431–1443. https://doi.org/10.1016/j.neucom.2008.06.011.

[31] Sainburg T, McInnes L, Gentner TQ.: Parametric UMAP embeddings for representation and semi-supervised learning. Available from: https://arxiv.org/abs/2009.12981.

[32] Santos-García I, Bascuñana P, Brackhan M, Villa M, Eiriz I, Brüning T, et al. The ABC transporter A7 modulates neuroinflammation via NLRP3 inflammasome in Alzheimer's disease mice [Journal Article]. Alzheimer's Research & Therapy.

2025;17(1):30. https://doi.org/10.1186/s13195-025-01673-2.

[33] Schwaiger-Haber M, Stancliffe E, Anbukumar DS, Sells B, Yi J, Cho K, et al. Using mass spectrometry imaging to map fluxes quantitatively in the tumor ecosystem [Journal Article]. Nat Commun. 2023;14(1):2876. https://doi.org/10.1038/s41467-023-38403-x.

[34] Ge J, Koutarapu S, Jha D, Dulewicz M, Zetterberg H, Blennow K, et al. Tetramodal Chemical Imaging Delineates the Lipid-Amyloid Peptide Interplay at Single Plaques in Transgenic Alzheimer's Disease Models [Journal Article]. Anal Chem. 2023;95(10):4692–4702. https://doi.org/10.1021/acs.analchem.2c05302.

[35] Xie YR, Castro DC, Rubakhin SS, Trinklein TJ, Sweedler JV, Lam F. Multiscale biochemical mapping of the brain through deep-learning-enhanced high-throughput mass spectrometry. Nat Methods. 2024;21(3):521–530. https://doi.org/10.1038/s41592-024-02171-3.

[36] Enzlein T, Lashley T, Sammour DA, Hopf C, Chavez-Gutierrez L. Integrative Single-Plaque Analysis Reveals Signature Abeta and Lipid Profiles in the Alzheimer's Brain [Journal Article]. Anal Chem. 2024;96(24):9799–9807. https://doi.org/10.1021/acs.analchem.3c05557.

[37] Mendis LH, Grey AC, Faull RL, Curtis MA. Hippocampal lipid differences in Alzheimer's disease: a human brain study using matrix-assisted laser desorption/ionization-imaging mass spectrometry [Journal Article]. Brain Behav. 2016;6(10):e00517. https://doi.org/10.1002/brb3.517.

[38] Veloso A, Fernandez R, Astigarraga E, Barreda-Gomez G, Manuel I, Giralt MT, et al. Distribution of lipids in human brain [Journal Article]. Anal Bioanal Chem. 2011;401(1):89–101. https://doi.org/10.1007/s00216-011-4882-x.

[39] Brackhan M, Calza G, Lundgren K, Bascunana P, Bruning T, Soliymani R, et al. Isotope-labeled amyloid-beta does not transmit to the brain in a prion-like manner after peripheral administration [Journal Article]. EMBO Rep. 2022;23(7):e54405. https://doi.org/10.15252/embr.202154405.

[40] Bascunana P, Brackhan M, Pahnke J. Machine Learning-Supported Analyses Improve Quantitative Histological Assessments of Amyloid-beta Deposits and Activated Microglia [Journal Article]. J Alzheimers Dis. 2021;79(2):597–605. https://doi.org/10.3233/JAD-201120.

[41] Radde R, Bolmont T, Kaeser SA, Coomaraswamy J, Lindau D, Stoltze L, et al. Abeta42-driven cerebral amyloidosis in transgenic mice reveals early and robust pathology [Journal Article]. EMBO Rep. 2006;7(9):940–6. https://doi.org/10.1038/sj.embor.7400784.

[42] Andersen MK, Hoiem TS, Claes BSR, Balluff B, Martin-Lorenzo M, Richardsen E, et al. Spatial differentiation of metabolism in prostate cancer tissue by MALDI-TOF MSI [Journal Article]. Cancer Metab. 2021;9(1):9. https://doi.org/10.1186/s40170-021-00242-z.

[43] Kibbe RR, Muddiman DC. Quantitative mass spectrometry imaging (qMSI): A tutorial. Journal of Mass Spectrometry. 2024;59(4):e5009. https://doi.org/10.1002/jms.5009.

[44] Winsor CP. Robust Estimation and Winsorization. Bulletin of the Harvard University Public Health. 1941;23:179–184.

[45] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. Advances in Neural Information Processing Systems. 2019;32. https://doi.org/10.48550/arXiv.1912.01703.

[46] Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:14126980. 2014;https://doi.org/10.48550/arXiv.1412.6980.

[47] Sunkin SM, Ng L, Lau C, Dolbeare T, Gilbert TL, Thompson CL, et al. Allen Brain Atlas: an integrated spatio-temporal portal for exploring the central nervous system [Journal Article]. Nucleic Acids Res. 2013;41(Database issue):D996–D1008. https://doi.org/10.1093/nar/gks1042.

4.10 Extended data

4.10.1 More visualization examples

Figure 16 shows an overview of UMAP, *SALO*, *Spear*, *PR3D*, and *TOP3* of one hemisphere of our mouse brain dataset (Table 1). The images show a coronary cut mouse brain hemisphere at bregma position –3.08 mm of a C57BL/6J mouse (commonly used B6 genetic background strain) and a schematic representation of the same location in the Allen mouse brain atlas [47]. UMAP reveals some gross details in the brain stem, e.g. the thalamus (TH) regional differences, the substantia nigra (SN), but also details in the hippocampus formation (HPF, green region in the Allen atlas), such as dentate gyrus (DG neurons) and to a lesser extent Cornu Ammonis field 3 neurons (CA3). Regarding the hippocampus formation (HPF), SALO, Spear, and PR3D reveal more details than UMAP and TOP3 with PR3D showing contrasted neuronal bands (CA1 and CA3, DG, all in dark pinkish) versus connections between them (layered pinkish or greenish). Spear especially highlights the molecular layer of the DG (DG-mo). PR3D highlights the layering and different extent of myelinization of the cortex (CTX, degree of myelinization as extent of the green signal). TOP3 shows a clear representation of the neuronal bands of the hippocampus formation (red, CA1, CA3, and DG) versus an orange background.

Figure 17 shows a zoom-in cross section of the hippocampal structures revealing different levels of details. UMAP often shows a homogeneous area and misses very distinct anatomical features (arrow 1 - basal layer of the dentate gyrus, arrow 2 - neurons in CA1) that are visible in the new visualizations. Arrow 3 points to single cells detected in the dentate gyrus isthmus using *PR3D*.

Fig. 16 Overview of UMAP, *SALO*, *Spear*, *PR3D*, and *TOP3* visualizations revealing different levels of anatomical details, e.g. in thalamus (TH), Substantia nigra (SN), hippocampus, and brain stem related to global data structure preservation.

Hippocampus

Fig. 17 Hippocampal details in UMAP, *SALO*, *Spear*, *PR3D*, and *TOP3* visualizations revealing different levels of details related to global data structure preservation.

4.10.2 Visualization benchmarking results with more metrics

Here we share more benchmarking results, with custom metrics that not have been used before for MSI. We share these here for additional future comparison as it may be valuable to collect more metrics for evaluation, but do not place them in the main result section to avoid confusion about our method performance with the established metrics used before for MSI. Figure 18 shows the benchmark results for our custom Saliency metric. Figure 19 shows the benchmark result for the Spearman correlation

metric. Our *SALO* is the best method in all datasets. *TOP3* and PR don't perform well on the Metaspace dataset under the Cosine distance.

Fig. 18 Benchmarking results for MSI visualizations the Saliency metric.

Fig. 19 Benchmarking results for MSI visualizations the Spearman rank correlation metric.

4.11 Supplementary information

Not applicable.

4.12 Acknowledgements

We thank Dr. Jorunn Stamnes for performing MSI measurements of the mouse brains and technical discussions, and Thomas Brüning for cutting and preparing the mouse brain hemispheres. Mass spectrometry-based analyses of mouse brains were performed at the Proteomics Core Facility at the University of Oslo/Oslo University Hospital, which is supported by the Core Facilities Program of the South-Eastern Norway Regional Health Authority (HSØ) and NAPI (www.napi.uio.no, NFR, Norway; 295910).

4.13 Funding

J.P. received funding from Norges forskningsråd (NFR, Norway; 327571 (PETABC), 295910 (NAPI, www.napi.uio.no)), HelseSØ (Norway; 2022046), and the EIC Pathfinder Open Challenges program (EU commission; 101185769).

4.14 Conflict of interest

The authors do not declare competing interests.

4.15 Ethics approval and consent to participate

Not applicable.

4.16 Consent for publication

Not applicable.

4.17 Author contribution

J.G. analyzed the MSI data and wrote and revised the manuscript. J.P. analyzed the MSI data, wrote and revised the manuscript, acquired funding, and provided infrastructure.