Technical explanation of the formulae in ”Truthful visualizations for mass spectrometry imaging enable high spatial resolution interactive m/z mapping and exploration”

Jacob Gildenblat and Jens Pahnke

Purpose and notation

This supplementary note explains the mathematical notation and formulae used in the main manuscript. The goal is to make the visualization objectives and evaluation metrics explicit while keeping the description connected to the practical MSI workflow.

An MSI dataset is represented as a three-dimensional array

$$rows \times cols \times N,$$

where $rows$ and $cols$ are the two spatial image dimensions and N is the number of measured *m/z* features. After flattening the spatial dimensions, each tissue pixel is treated as one spectrum

$$x_i = (x_{i1}, x_{i2}, \ldots, x_{iN}) \in \mathbb{R}^N,$$

where i indexes pixels and x_{ij} is the intensity of the j-th *m/z* bin in pixel i. A visualization maps each high-dimensional spectrum x_i to a low-dimensional display coordinate or color vector

$$y_i \in \mathbb{R}^2 \quad \text{or} \quad y_i \in \mathbb{R}^3.$$

For RGB-like visualizations, y_i is interpreted as a color assigned to the original spatial pixel.

Scheme 1: MSI image cube, spectra, and molecular-color map

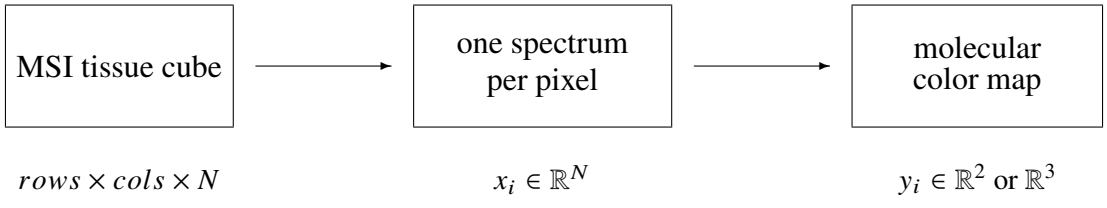

This scheme summarizes the notation in MSI terms: each tissue pixel contains a full mass spectrum, and the visualization converts spectral similarity and difference into a molecular color map at the original tissue position.

Distance functions in MSI feature space

The main text discusses complementary distance behaviors because different molecular changes are captured by different metrics.

The Euclidean distance between two spectra is

$$d_2(x_i, x_j) = \left(\sum_{k=1}^{N} (x_{ik} - x_{jk})^2 \right)^{1/2}.$$

This distance increases when many *m/z* features change together, even if each individual feature changes only moderately. It is therefore useful for coordinated molecular shifts across broad parts of the spectrum.

The cosine distance can be written as

$$d_{\cos}(x_i, x_j) = 1 - \frac{x_i \cdot x_j}{\|x_i\|_2 \|x_j\|_2}.$$

It compares spectral direction rather than absolute magnitude and is useful when relative spectral composition is more important than total intensity.

The L_∞ distance is

$$d_\infty(x_i, x_j) = \max_{1 \le k \le N} |x_{ik} - x_{jk}|.$$

This distance is dominated by the strongest single *m/z* difference. It is useful for hot-spot-like changes, such as a small number of ions that strongly distinguish one tissue structure from another.

Scheme 2: MSI interpretation of distance metrics

Distance	MSI signal pattern emphasized	Example in tissue interpretation
Euclidean	many ions shift together	gray–white matter or cortex–medulla gradients
Cosine	relative spectral composition changes	similar total intensity but different lipid profile
L_∞	one ion changes strongly	plaque-, neuron-, or deposit-associated ion hot spot

The methods use these complementary views because MSI tissue contrast may arise from broad lipid-program differences or from a small number of highly localized diagnostic ions.

Ranked distances and reference points

SALO and SPEAR preserve global structure by comparing ranked distances from representative reference points. Let r denote one selected reference pixel. For all pixels i, the high-dimensional distance from the reference point is

$$D^X_{r,i} = d(x_r, x_i),$$

and the visualization-space distance is

$$D^Y_{r,i} = d(y_r, y_i).$$

The rank vectors are

$$R^X_{r,i} = \mathrm{rank}(D^X_{r,i}), \qquad R^Y_{r,i} = \mathrm{rank}(D^Y_{r,i}).$$

Ranks describe ordering rather than absolute distance values. If two pixels are far apart from the reference point in MSI feature space, they should also be far apart in the visualization. Conversely, if they are similar in MSI feature space, they should remain close or visually related after projection.

The main manuscript notes that SALO and SPEAR use a differentiable approximation of the ranking operation. This is necessary because ordinary ranks are discrete and cannot be optimized directly by gradient-based methods. A smooth rank approximation replaces the hard sorting operation with a differentiable surrogate, allowing the visualization colors or coordinates to be updated during optimization.

Scheme 3: preserving distances from an MSI reference pixel

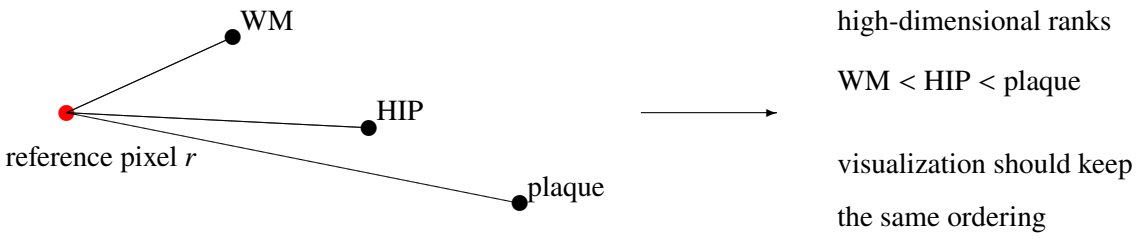

For each reference pixel, other tissue pixels are ordered by spectral distance. SALO and SPEAR optimize the molecular-color map so that, for example, hippocampal, white-matter, plaque-associated, or inflammatory pixels do not collapse into visually indistinguishable colors when their spectra are different.

Fusion of Euclidean and L_∞ ranks

To combine coordinated spectral changes with single-ion-dominated changes, the manuscript describes fusing Euclidean and L_∞ information by taking the maximum of their ranks. For a reference point r and pixel i, this can be written as

$$R^X_{r,i} = \max\left(R^{X,2}_{r,i}, R^{X,\infty}_{r,i}\right),$$

where $R^{X,2}_{r,i}$ is the rank obtained from Euclidean distances and $R^{X,\infty}_{r,i}$ is the rank obtained from L_∞ distances. This rule treats a pair of pixels as globally separated if either the distributed spectral pattern or a strong individual *m/z* difference indicates separation.

Scheme 4: combining broad molecular shifts and ion hot spots

MSI comparison	Euclidean rank	L_∞ rank	fused rank
broad lipid-region shift	8	2	$\max(8, 2) = 8$
localized ion hot spot	3	9	$\max(3, 9) = 9$
mixed spectral difference	5	5	$\max(5, 5) = 5$

A high fused rank means that two tissue pixels should remain visibly separated if either a broad spectral program or a single localized ion identifies them as molecularly different.

SPEAR objective

SPEAR optimizes Spearman rank correlation between high-dimensional ranked distances and visualization-space ranked distances. For one reference point r, the Spearman correlation can be expressed as the Pearson correlation of the two rank vectors:

$$\rho_r = \frac{\sum_i (R^X_{r,i} - \bar{R}^X_r)(R^Y_{r,i} - \bar{R}^Y_r)}{\sqrt{\sum_i (R^X_{r,i} - \bar{R}^X_r)^2}\sqrt{\sum_i (R^Y_{r,i} - \bar{R}^Y_r)^2}}.$$

Here, $\bar{R}_r^X$ and $\bar{R}_r^Y$ are the mean ranks over all pixels for reference point r. SPEAR maximizes the average correlation over reference points:

$$\mathcal{L}_{\text{SPEAR}} = -\frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \rho_r,$$

where $\mathcal{R}$ is the set of selected reference points. The negative sign indicates that the optimization minimizes a loss while effectively maximizing rank correlation.

Scheme 5: SPEAR preserves the order of tissue similarities

MSI spectra ranked from a reference pixel		visualization should show the same order
reference → WM → HIP → plaque	$\Longrightarrow$	reference → WM → HIP → plaque
small spectral distance		small color distance
large spectral distance		large color distance

SPEAR rewards a visualization when the tissue pixels keep the same rank order around each reference pixel: pixels with similar spectra remain visually close, whereas molecularly distinct pixels remain visually separated.

SALO saliency objective

SALO uses the same ranked-distance principle but applies a saliency-based objective. The aim is not only to make rank orderings correlated, but to prevent visually important MSI differences from collapsing in the display. The manuscript describes this as requiring visualization-space distance ranks to be at least as high as MSI-space distance ranks.

A simplified form of this condition is

$$R_{r,i}^Y \geq R_{r,i}^X.$$

Violations occur when

$$R_{r,i}^Y < R_{r,i}^X,$$

meaning that two pixels are relatively distant in molecular space but appear too close in the visualization. A hinge-style loss for this idea is

$$\mathcal{L}_{\text{SALO}} = \frac{1}{|\mathcal{R}|M} \sum_{r \in \mathcal{R}} \sum_i \max\left(0, R_{r,i}^X - R_{r,i}^Y\right),$$

where M is the number of compared pixels per reference point. In practice, SALO uses differentiable rank approximations and optimizes colors in a perceptually meaningful color representation. The main text specifies LAB color space because equal numeric changes in LAB more closely approximate equal perceived color changes than equal numeric changes in raw RGB.

Scheme 6: SALO prevents molecularly distinct tissue pixels from collapsing into the same color

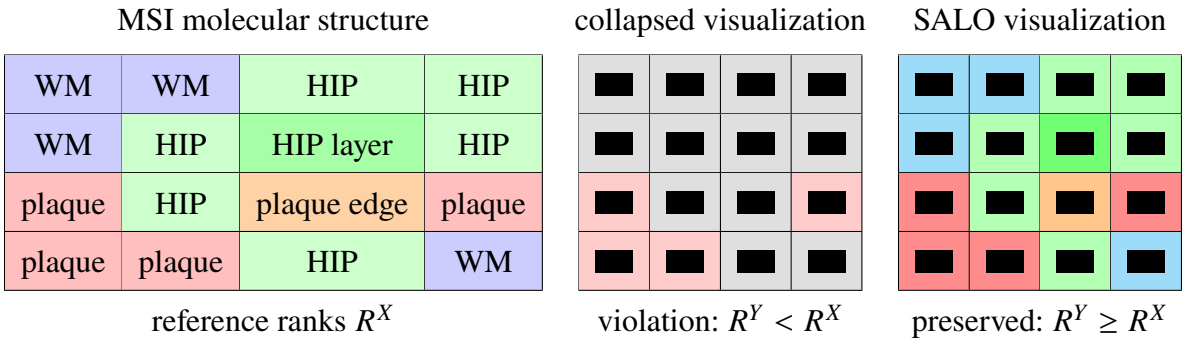

SALO penalizes the middle case, where molecularly distinct MSI pixels are assigned nearly the same visual color. The optimized map on the right keeps white matter, hippocampal sublayers, plaque edges, and plaque cores visually separated when their spectra are separated in the MSI data.

Visualization-space distances and LAB color

For color visualizations, a pixel color can be represented as

$$y_i = (L_i, a_i, b_i),$$

where L is lightness and a, b are opponent color channels in LAB color space. A simple visualization-space distance is

$$d_{LAB}(y_i, y_j) = \left((L_i - L_j)^2 + (a_i - a_j)^2 + (b_i - b_j)^2\right)^{1/2}.$$

Using LAB distances helps the objective align numerical color separation with human-visible contrast. This is important because the output is interpreted visually by researchers and pathologists.

TOP3 mapping

TOP3 is a lightweight visualization that maps dominant ion intensities to color channels. For each pixel i, let $t_1(i), t_2(i), t_3(i)$ be the indices of the three strongest selected or normalized *m/z* signals. A simplified TOP3 color assignment is

$$y_i = \left(s_{i,t_1(i)}, s_{i,t_2(i)}, s_{i,t_3(i)}\right),$$

where s_{ij} is a scaled intensity for *m/z* feature j in pixel i. The result is a continuous RGB-like representation in which the most prominent local molecular signals define the color. Unlike clustering, TOP3 does not force pixels into discrete classes.

Scheme 7: TOP3 as a dominant-ion color mixture

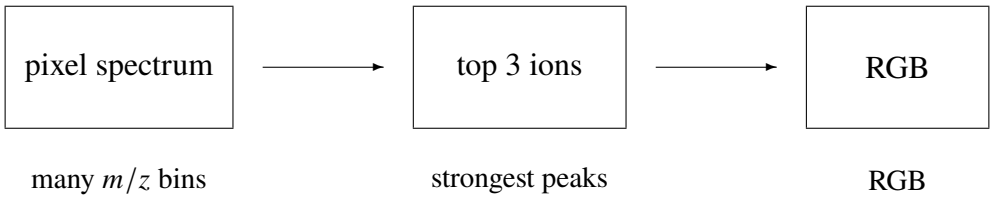

In MSI terms, TOP3 asks which three ion signals dominate a pixel after scaling and uses them as a local color mixture. This can highlight tissue structures with strong characteristic peaks while remaining computationally lightweight.

PR3D percentile mapping

PR3D uses percentile-based scaling to convert intensity distributions into a three-channel visualization. If $p_{low,j}$ and $p_{high,j}$ are lower and upper percentile cutoffs for feature or channel j, a scaled intensity can be written as

$$s_{ij} = \mathrm{clip}\left(\frac{x_{ij} - p_{low,j}}{p_{high,j} - p_{low,j}}, 0, 1\right).$$

The clip function truncates values below 0 to 0 and values above 1 to 1. Percentile scaling suppresses extreme outliers and enhances contrast in the intensity range most relevant for visualization. The main manuscript notes that PR3D is sensitive to percentile settings; this follows directly from

the formula because changing $p_{low,j}$ or $p_{high,j}$ changes which intensity differences are visually emphasized.

Scheme 8: PR3D contrast window for MSI ion intensities

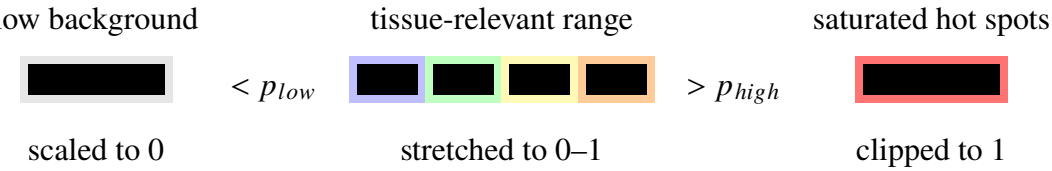

PR3D chooses a percentile window for ion intensities. Background-like values below the window are suppressed, the tissue-relevant middle range is expanded into visible contrast, and extreme hot spots are clipped so they do not dominate the entire molecular image.

Benchmark correlations

Benchmarking compares whether distances in the visualization preserve distances in the MSI feature space. A common global preservation score is Pearson correlation between high-dimensional distances and visualization-space distances:

$$r = \frac{\sum_q (D_q^X - \bar{D}^X)(D_q^Y - \bar{D}^Y)}{\sqrt{\sum_q (D_q^X - \bar{D}^X)^2}\sqrt{\sum_q (D_q^Y - \bar{D}^Y)^2}}.$$

Here, q indexes sampled pixel pairs or reference-point–pixel pairs. A higher value means that molecularly distant pixels tend to remain visually distant and molecularly similar pixels tend to remain visually similar.

The Spearman benchmark uses the same correlation formula after replacing distances by ranks. It therefore evaluates whether the ordering of distances is preserved, even if the exact numeric distances are distorted.

ROI-to-*m/z* scoring

The manuscript describes selecting regions of interest and identifying *m/z* values enriched in one region compared with another. For two ROIs, A and B, the mean intensity of feature j in each

region is

$$\mu_{A,j} = \frac{1}{|A|} \sum_{i \in A} x_{ij}, \qquad \mu_{B,j} = \frac{1}{|B|} \sum_{i \in B} x_{ij}.$$

A simple enrichment score is the difference in means

$$\Delta_j = \mu_{A,j} - \mu_{B,j},$$

or a fold-change-like ratio

$$F_j = \frac{\mu_{A,j} + \epsilon}{\mu_{B,j} + \epsilon},$$

where ϵ is a small constant used to avoid division by zero. Statistical comparison can additionally use a two-sample statistic such as

$$t_j = \frac{\mu_{A,j} - \mu_{B,j}}{\sqrt{s^2_{A,j}/|A| + s^2_{B,j}/|B|}},$$

where $s^2_{A,j}$ and $s^2_{B,j}$ are within-ROI variances. High-scoring *m/z* values are then mapped back into image space as ion images or red/green overlays.

Scheme 10: pathologist-guided ROI-to-*m/z* mapping

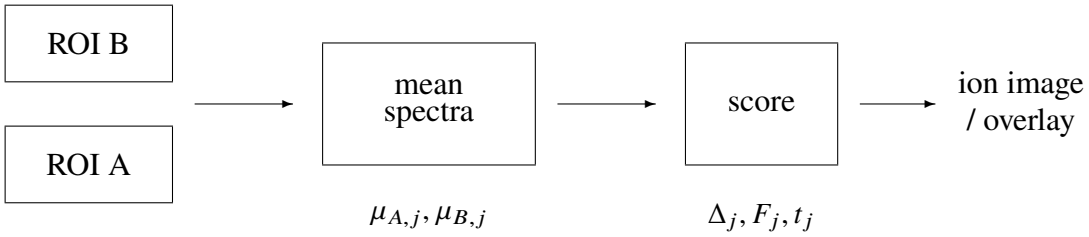

The ROI formulae therefore connect pathologist- or researcher-selected tissue regions to ranked candidate ions and then back to ion images or red/green overlays for spatial confirmation in the same MSI dataset.

Binning precision

The supplementary binning panels describe bins per Dalton. If B bins are used per Dalton, then the nominal bin width is

$$\Delta m = \frac{1}{B} \text{ Da}.$$

Thus 5, 10, 20, and 40 bins per Dalton correspond to 0.2, 0.1, 0.05, and 0.025 Da, respectively. Expressed in millidaltons, these are 200, 100, 50, and 25 mDa. Increasing B improves nominal *m/z* precision but increases memory and computational cost because more spectral features must be stored and processed.

Scheme 9: binning precision and computational cost

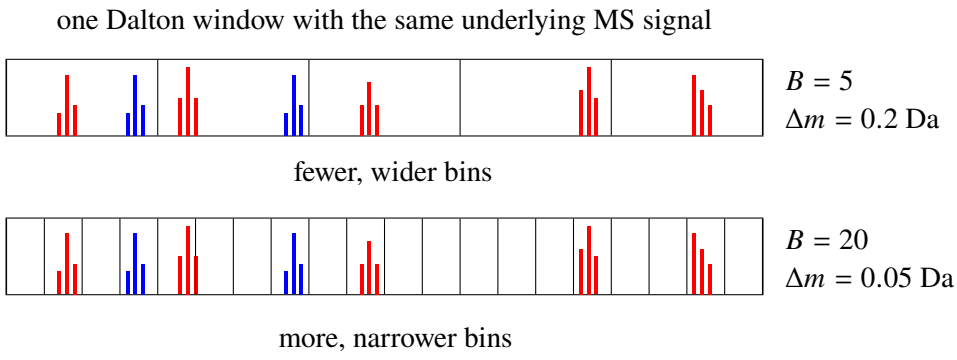

This scheme illustrates the formula $\Delta m = 1/B$: increasing the number of bins per Dalton narrows each *m/z* bin and improves nominal mass precision, but it also increases the number of stored features and therefore memory and runtime.

Interpretation

Together, these formulae formalize the central idea of the manuscript: MSI visualizations should preserve molecular relationships that are relevant for anatomical and pathological interpretation. SPEAR does this by maximizing rank correlation, SALO does this by penalizing visual collapse of molecularly separated pixels, TOP3 and PR3D provide fast continuous intensity-based color mappings, and ROI-to-*m/z* scoring connects visual regions back to molecular features.