A Novel Peak Shape-Aware Approach for Mass Alignment in Mass Spectrometry

Thomas Vanhemel[1][a], Melanie Nijs[1,2], Angeliki Birmpili[3], Tim Hendriks[3], Eva Cuypers[3], Bart De Moor[1,4]

[1]*KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics*

[2]*University of Antwerp, Department of Communication Sciences*

[3]*Maastricht University, The Maastricht MultiModal Molecular Imaging (M4I) Institute, Division of Imaging Mass Spectrometry (IMS)*

[4]*Fellow IEEE, SIAM, IFAC*

Keywords: Mass Spectrometry, Mass Spectrometry Imaging, Mass Shift, Label-Free Mass Alignment

Abstract: **RATIONALE** During mass spectrometry measurements, mass shifts may be inadvertently introduced due to instrumental drift and calibration inaccuracies, potentially compromising the accuracy of subsequent data analysis. This work presents a novel, label-free algorithm to improve relative mass alignment between mass spectra. The warping function is modeled as a natural cubic spline, a suitable model for gradual, nonlinear mass shifts.

METHODS The algorithm is validated on a dataset generated from human glioblastoma multiforme samples using matrix-assisted laser desorption/ionization time-of-flight mass spectrometry imaging (MALDI-TOF MSI) and rapid evaporative ionization mass spectrometry (REIMS).

RESULTS The algorithm considerably reduces the mass dispersion of the dataset and improves the similarity to a reference mass spectrum for MALDI-MSI and REIMS data. It is demonstrated that the proposed method reliably corrects for severe mass shifts. The algorithm is competitive in speed and mass dispersion reduction compared to MSIWarp, a common method to correct for mass misalignment.

CONCLUSIONS This paper presents a novel algorithm to reduce relative mass misalignment for mass spectrometry, validated on an extensive dataset. Thanks to the use of profile data, the peak shapes contribute to the computation of the warping function, and the warping function is approximated in a robust manner. An open-source Python implementation of the proposed methodology will be made accessible on GitHub: https://github.com/VanhemelThomas/psalign.

1 INTRODUCTION

Mass spectrometry is an analytical technique used to measure the abundance of ionized molecules in a sample (Gross, 2017). This technique results in a mass spectrum consisting of the ion count per mass-to-charge (m/z) ratio. Mass spectrometry imaging (MSI) measures mass spectra in a spatially resolved and systematic manner (Caprioli et al., 1997). The resulting high dimensional image is a rich source of spatial and chemical information. The high mass spectral resolution is essential for correctly identifying compounds present in the sample. However, the measured mass spectra have empirically been observed to suffer from a shift in m/z values across the sample. This shift is typically more prevalent for instruments containing time-of-flight (TOF) detectors, but has also been observed for other instruments (Ràfols et al., 2018). Mass misalignment can occur for a multitude of reasons in TOF instruments, for instance, thermal distortion of the TOF tube (Tracy et al., 2008), contamination of the ion source (Eriksson et al., 2020), inhomogeneous matrix deposition (McCann et al., 2021), topography of the tissue and variations in sample thickness (McCann et al., 2021) (Norris et al., 2007) can contribute to the mass shift phenomenon. This phenomenon can be addressed at the point of measurement or retrospectively through data-driven correction methods. In this work, the former is referred to as *mass calibration* and the latter as *mass alignment*.

To resolve misalignment, mass calibration can be performed via internal and external mass calibration (Gross, 2017). During internal mass calibration,

[a] https://orcid.org/0009-0003-3306-2101

a calibration compound can be introduced to a second inlet of the mass spectrometer, or the compound can be mixed with the analyte prior to analysis. This could potentially modify the analyte, hence the former is preferred. During external mass calibration, a calibration compound with known mass spectrum is fed to the mass spectrometer. Now, the mass spectrum is calibrated by determining a warping function from the measured spectrum and the calibration peaks. The warping function maps measured m/z values to those of the calibration peaks. Figure 1 depicts a toy example of such a warping function. This warping function is applied to all subsequent measurements. Hence, this calibration technique is unable to correct for mass shifts that arise due to changing conditions during the experiment. Gobom et al. (Gobom et al., 2002) observed the mass shift phenomenon after external calibration and noted that the mass shift changes over time. Therefore, internal mass calibration obtains more accurate calibration results. However, the application of the calibration compound across the tissue sample in MSI could have undesirable ion suppression effects. Thus, internal calibration is not ubiquitous in an MSI setting (Boskamp et al., 2020).

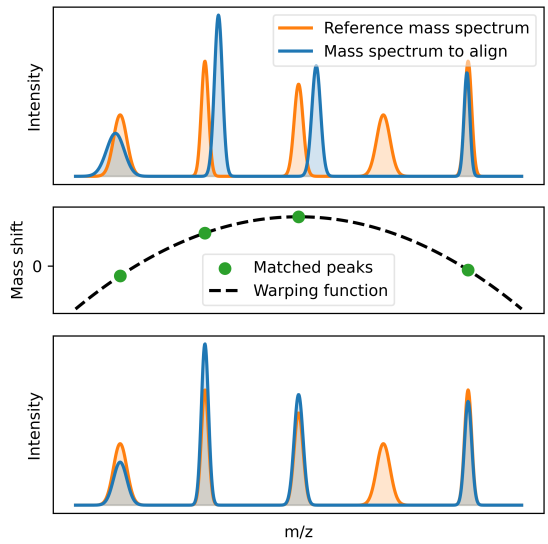

Figure 1: A toy example of a warping function: (top) a mass spectrum with mass misalignment with respect to the reference mass spectrum; (middle) the warping function which alters the m/z values corresponding to the mass spectrum to align, such that alignment between the spectra improves; (bottom) aligned mass spectra after warping the mass spectrum.

1.1 Mass alignment

Computational mass alignment approaches warp all spectra to a common reference spectrum, which can be constructed using known compounds in the sample or, for instance, the average mass spectra of a mass spectrometry image. These data-driven alignment approaches typically reduce the dimensionality of the problem by performing peak picking and subsequently matching the peaks. Later, the spurious matches are identified and removed before the final warping function is computed. This paper proposes an alternative method which does not discard the shapes of the peaks or parts of the signal with lower intensity, and is therefore able to reliably correct large mass shifts.

Most existing algorithms perform peak matching by extracting peaks from the profile data, identifying peaks in the neighborhood of a given peak and considering the peaks matched if the m/z difference between the peaks is below a predefined threshold (Bocker and Makinen, 2008)(Kulkarni et al., 2015)(Boskamp et al., 2020)(Denker et al., 2024)(Eriksson et al., 2020)(La Rocca et al., 2021). These algorithms differ in the manner spurious matches are identified and in the manner the final warping is computed. Böcker and Mäkinen (Bocker and Makinen, 2008) proposed several methods for the detection of spurious matches, all of these compute a linear function to map between peaks using a combinatorial algorithm robust to outliers. Spurious matches are eliminated based on the distance to the linear function and the final warping is computed by fitting a polynomial function (up to order 2) to the true matches. Kulkarni et al. (Kulkarni et al., 2015) expanded on this method by ordering the pixel pairs according to the similarity between their peaks. A reference spectrum is iteratively generated based on the pixel pair ordering.

Boskamp et al. (Boskamp et al., 2020) improve the alignment by exploiting the fact that tryptic peptides are observed to be clustered around equidistant masses across the m/z range. These peaks are matched with the mass spectra, resulting in a Kendrick mass defect map that encodes the difference between the measured masses and the nearest, exact peptide masses. A statistical analysis of the map provides local warping functions and the complete warping function is obtained by linearly interpolating the remaining m/z bins. The method is not universally applicable due to the dependence on the chemical noise background, and yields erroneous results in spectra dominated by matrix peaks and located in off-tissue areas. Recently, Denker et al. (Denker et al., 2024) proposed a neural network to improve the work of

Boskamp et al. The proposed neural network maps the Kendrick mass defect map to a warping function. Hence, the efficacy of this method is contingent on the presence of the chemical noise background. Additionally, the proposed method only considers mass defects within the range of roughly 0.5 *Da* due to the circular definition of the Kendrick mass defect map.

Eriksson et al. (Eriksson et al., 2020) developed another label-free approach, called MSIWarp, which represents the peaks in the mass spectra as Gaussians and eliminates false matches by performing a random sample consensus (RANSAC) procedure (Fischler and Bolles, 1981). However, the approximation of peaks using Gaussian functions can lead to inaccuracies in case of peak asymmetry, a characteristic frequently observed in proteomics and metabolomics studies (Wijetunge et al., 2015). This approach models the warping function as a piecewise linear function that maximizes a similarity metric, which is based on the overlap of the Gaussians between spectra. The warping function is computed with dynamic programming.

In the work of La Rocca et al. (La Rocca et al., 2021) a Gaussian kernel is used to estimate the probability density function of the mass errors. Subsequently, matches are discarded based on the distance to the maximum of the distribution and RANSAC regression is used since the matches are not necessarily outlier-free.

An example of an intrasample alignment method for which peak picking is not required is the work of Ràfols et al. (Ràfols et al., 2018). Here, an initial linear warping is performed based on the cross-correlation of the upper and lower part of the mass spectrum and the reference spectrum. Afterwards, nonlinear Loess smoothing is applied to predict the calibrated m/z axis based on the relationship between reference masses and the average spectrum. Although there is a nonlinear step, the intrasample nonlinear relationships are not corrected.

This paper proposes a novel, peak-picking-free approach for computing optimal nonlinear warping functions in mass alignment. The method directly estimates a smooth cubic spline warping function that maximizes the similarity between a reference mass spectrum and a warped mass spectrum. By avoiding the potentially unreliable peak-picking step (Abdelmoula et al., 2021), this approach preserves the complete spectral signal, thereby enhancing robustness to noise. The robustness of the method is demonstrated by aligning spectra exhibiting significant mass shifts and correcting the mass shift between spectra across different MS modalities: matrix-assisted laser desorption/ionization mass spectrometry imaging (MALDI-MSI) and rapid evaporative ionization mass spectrometry (REIMS).

2 METHODS

The goal of the alignment procedure is to compute a warping function that maximizes a similarity metric between a reference spectrum and a moving spectrum. This procedure is applied to all spectra to align them to a common reference mass spectrum. Evidently, the efficacy of data-driven mass alignment depends on the availability of sufficient shared information between the reference and moving spectrum. The methodology for mass alignment presented in this paper consists of four steps. First, the mass spectrum is divided into segments of equal size and the maximum peak within that segment is identified. Second, a local window around that peak is compared with the reference spectrum, the shift that maximizes the similarity estimates the local mass shift. Third, a warping function is calculated based on the local mass shifts after an optional outlier detection step. Fourth, the parameters of the warping function are optimized to maximize a similarity metric. Figure 2 gives an overview of the proposed algorithm.

2.1 Mass alignment as a mathematical optimization problem

The warping function is modeled as a natural cubic spline function S, which is a piecewise cubic function that guarantees smoothness at the boundaries and zero second derivative at both ends (Boor, 1978). Therefore, splines are ideal for modeling gradual, nonlinear mass shifts. The spline function is uniquely defined by the knots (a_i, b_i) with $\mathbf{a} = (a_1, \ldots, a_M)$ and $\mathbf{b} = (b_1, \ldots, b_M)$ and can be evaluated at the point x with $\Phi(S(\mathbf{a}, \mathbf{b}), x)$, for example $\Phi(S(\mathbf{a}, \mathbf{b}), a_1) = b_1$. Smets et al. (Smets et al., 2019) reported that the cosine similarity performs well as a distance or similarity metric between mass spectra. Hence, the cosine similarity is chosen as the similarity metric and is bounded between 0 and 1 for nonnegative data. Now, it is possible to formulate the mass alignment problem as a mathematical optimization problem.

$$\min_{\mathbf{b}} \quad \mathcal{L}(\mathbf{b}) = \frac{-r(\mathbf{z}) \cdot q(\Phi(S(\mathbf{a}, \mathbf{b}), \mathbf{z}))}{\|r(\mathbf{z})\| \|q(\Phi(S(\mathbf{a}, \mathbf{b}), \mathbf{z}))\|} \tag{1a}$$

$$\text{s. t.} \quad b_m^0 - \Delta \leq b_m \leq b_m^0 + \Delta, \quad m = 1, \ldots, M \tag{1b}$$

The reference mass spectrum r is evaluated at the centers of the m/z bins $\mathbf{z}$ with $\mathbf{z} = (z_1, \ldots, z_N)$. In

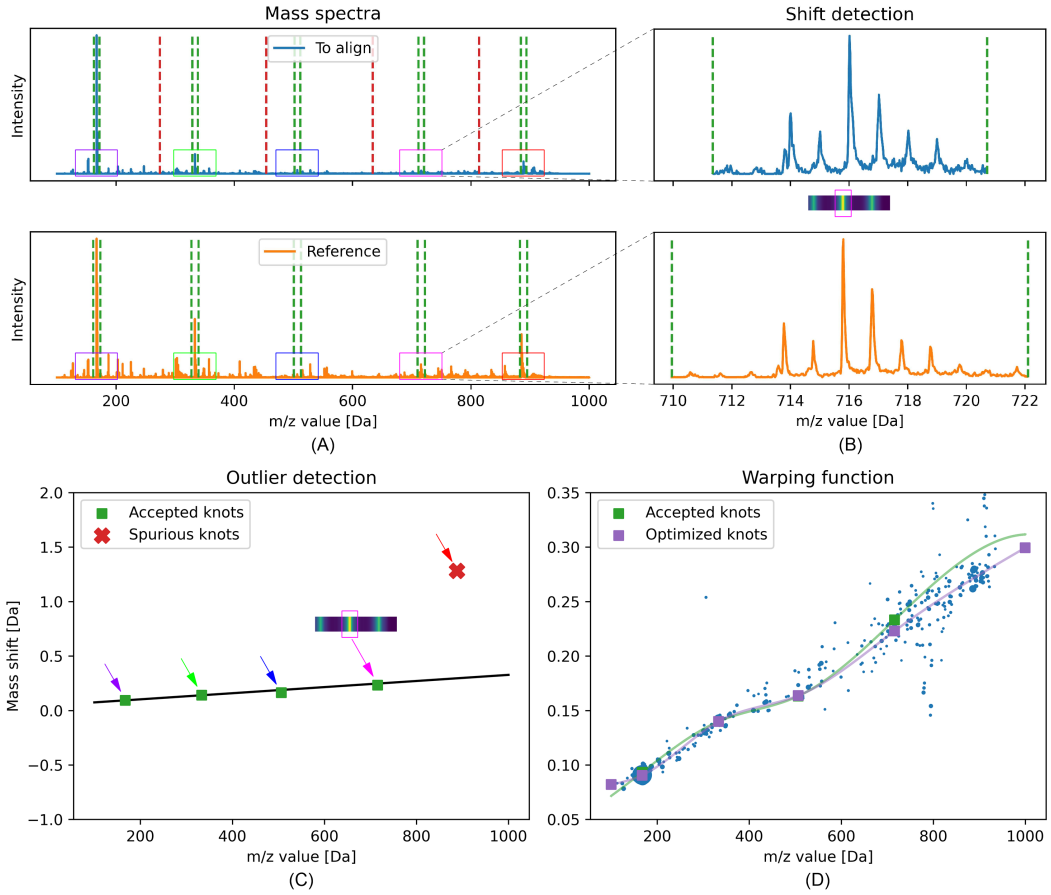

Figure 2: Overview of the proposed methodology. (A) The mass spectrum is divided into segments of equal size and a window around the largest peak is taken and a slightly larger window is extracted from the reference mass spectrum. (B) The mass spectrum window slides over the reference spectrum window in order to identify the optimal local mass shift. The similarity score across the shifts is visualized and the optimal local mass shift is highlighted. (C) An optional outlier detection procedure is applied to the detected local mass shifts or knots and outliers are removed. (D) The warping function is computed based on the knots and the knots are adjusted to maximize a similarity metric.

Equation (1a), q denotes the spectrum to be warped and is evaluated at m/z values shifted by the spline $S(\mathbf{a}, \mathbf{b})$. However, q will rarely be evaluated at the centers of the m/z bins due to continuous shifts applied by the warping function, thus it is linearly interpolated. Constraints (1b) impose that the spline cannot exceed a maximum mass shift Δ at the knots where $\mathbf{b}^0 = (b_1^0, \ldots, b_M^0)$ denotes the initial estimate of the optimization problem.

2.2 Obtaining a robust initial estimate

Given the abundance of peaks present in mass spectrometry, it is expected that the optimization problem has numerous suboptimal minimizers. Consequently, a global optimization algorithm would be advantageous; however, given the need to repeat the procedure for each pixel of a MS image, the associated computation time would rapidly become excessive. Hence, an alternative strategy is explored that uses a gradient-based algorithm to improve a good initial estimate.

In order to obtain a good initial estimate for the warping function, q is divided into segments, and the highest peak is found per segment, see Figure 2 (A). Then, a section containing the largest peak is extracted and matched to the reference spectrum, meaning that the local shift is found that maximizes the dot product between the reference spectrum and the section containing the largest peak, see Figure 2 (B). The use of the dot product facilitates the identification of corresponding peaks since the peak shape provides additional information for the determination of the mass shift, making the proposed method more robust. This information would be absent if the mass

spectrum were peak-picked or centroided. Optionally, an outlier detection approach is implemented to remove erroneous matches should they occur, see Figure 2 (C). A sliding window of size five is applied to consecutive detected mass shifts and their associated m/z values. Within each window, the mass shift trend is approximately linear, allowing the reliable detection of outliers using a relatively small window size. The detection of outliers is akin to the RANSAC (Fischler and Bolles, 1981) algorithm: (i) select two mass shifts and their associated m/z values from the window and fit a linear model; (ii) detect mass shifts of the window which are consistent with the linear model; (iii) repeat steps (i)-(ii) for all possible choices of two mass shifts out of the window of five and select the model which maximizes the number of mass shifts consistent with the model. If there are multiple solutions select the model with the smallest error on the consistent set of mass shifts. The mass shifts in (i) are not selected randomly, every possible combination of two mass shifts out of five is selected since there are only $C_2^5 = 10$ possible combinations. Consequently, the procedure is deterministic. As the window slides over all detected mass shifts, every mass shift receives an inlier or outlier vote each time it is part of the sliding window. Outliers are those mass shifts that receive more outlier votes than inlier votes and are ultimately discarded. The collection of accepted knots are expanded by knots at the boundaries of the m/z array and used for the initial estimate of the warping function. The addition of knots at the boundaries improves the control of the warping function in proximity to the boundaries, eliminating the need for extrapolation.

Once an initial estimate of the knots is obtained, it can be optimized by solving the optimization problem, see Figure 2 (D). At this point, the assumption is made that the initial estimate is close to the global optimum, hence the search space can be considerably reduced by opting for a small Δ value. L-BFGS-B (Zhu et al., 1997) is used to solve the optimization problem. It is a quasi-Newton method that updates an approximation of the inverse of the Hessian at every step and can handle simple box constraints on the input variables. It is empirically observed that the initial estimate considerably corrects the mass shift across the sample in most cases, and that solving the optimization problem is time-consuming. Therefore, performance metrics are also reported without solving the corresponding optimization problem for every spectra.

2.3 Evaluation metrics

Although ground truth with respect to mass alignment is generally unavailable, the performance of an alignment procedure can be estimated by examining the mass dispersion of the peaks and the similarity of the mass spectra to the average mass spectrum before and after alignment. The reference spectrum used is the average mass spectrum after Total Ion Current (TIC) normalization (Gross, 2017), unless explicitly indicated otherwise. In this text, this mass spectrum is referred to as the *average spectrum*. This spectrum usually contains shared information across pixels and consistent biological signal. To compute the mass dispersion of the data, the average mass spectrum is used to detect the most intense peaks. The parabolic centroiding algorithm (Robichaud et al., 2013) is employed to obtain an estimate of the corresponding peak mass. Then, for every mass spectrum in the sample, peaks are identified within a window around the peak mass, of width two times the peaks full width at half-maximum. The detected peak closest to the peak mass is chosen, and the standard deviation of the peak mass dispersion is computed.

A comprehensive evaluation of the methods requires the use of multiple metrics, as mass dispersion alone is insufficient to determine whether peaks originate from the same compound. While previous studies have primarily relied on mass dispersion to assess alignment quality, this work extends the evaluation by incorporating a similarity metric. This similarity metric provides stronger evidence for the correct correspondence of peaks. As described earlier, the cosine similarity is the chosen similarity metric.

3 MATERIALS

The developed method is evaluated through its application to a large dataset. The dataset is generated from stage IV primary human glioblastoma multiforme sections with a RapifleX MALDI Tissuetyper TOF mass spectrometer (Bruker Daltonics) in negative ionization mode (Hendriks et al., 2025). Specifically, the dataset includes 24 samples drawn from the larger cohort analyzed by Hendriks et al. (Hendriks et al., 2025). Furthermore, REIMS measurements were generated from the remaining tissue using an electrical surgical knife (Balog et al., 2010) connected to a Xevo G2-XS quadrupole time-of-flight mass spectrometer with a REIMS interface (Waters Corporation). Multiple incisions were made per sample. All patients ($>$ 11 years of age) were scheduled to undergo an elective neurosurgical procedure

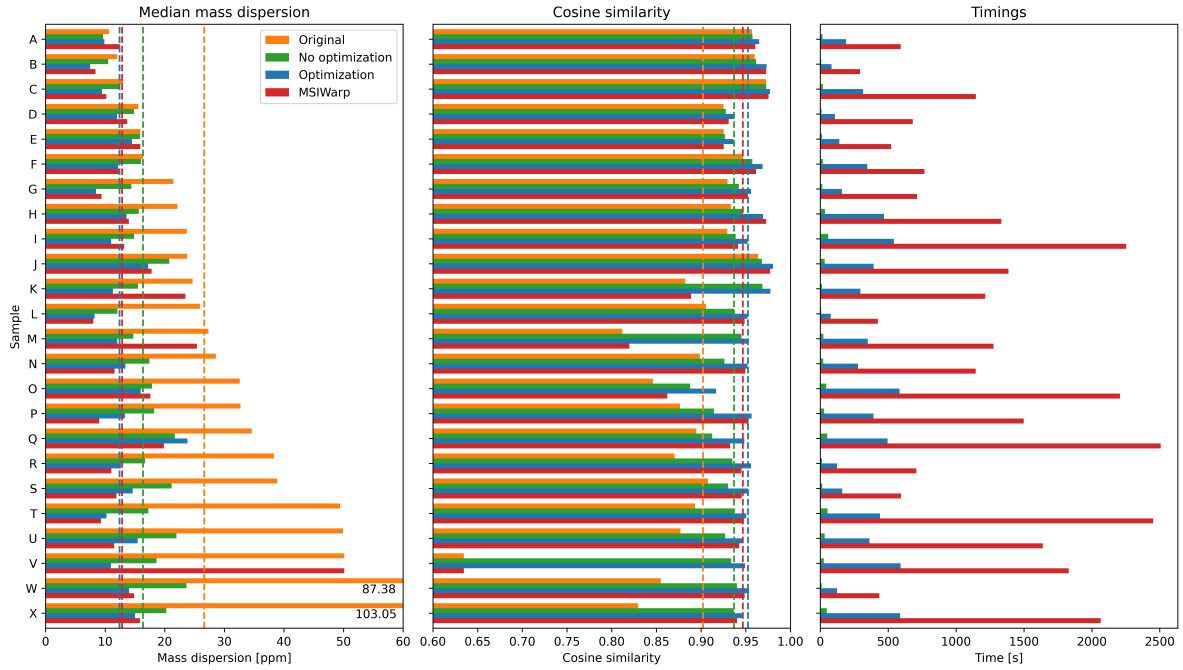

Figure 3: Comparison of evaluation metrics after applying our approach and MSIWarp to the dataset. Left: Median mass dispersion per sample of the 100 most dominant peaks of the average mass spectrum. The dashed line represents the median of the median mass dispersions. Middle: Mean cosine similarity of the mass spectra and the average mass spectrum after alignment, or before alignment for the original data. The dashed line represents the median of the mean cosine similarities. Right: Execution time of the algorithms per sample.

under total intravenous general anaesthesia. This single-centre, prospective observational study was performed at Leuven University hospitals, Gasthuisberg (Leuven, Belgium). The study protocol was approved by the Ethics Committee Research UZ/KU Leuven (S60291) and the project was registered under the UZ Leuven Tissue Bank. The m/z range considered in this paper is 100–1000 Da. Supplementary information on the dataset can be found in Table S1. No additional preprocessing was performed on the dataset.

After reducing the intrasample misalignment of the MALDI images, the average MALDI mass spectrum can be aligned to the average REIMS measurement of the sample. This demonstrates the robustness of the developed approach.

The proposed methodology necessitates an identical array of m/z values between the reference spectrum and the moving spectrum. In the case of profile data, the m/z array of the moving spectrum can be binned to the m/z array of the reference spectrum. For sparse centroided data, a shared m/z array can be constructed based on all the measured m/z values in subsequent measurements. The initial m/z array is chosen as the array of the measurement with the largest number of m/z bins. The m/z values are then added to the shared m/z array iteratively, provided that the distance from the m/z value to the shared m/z array is greater than the smallest distance between m/z values in a sample across the measurements. Finally, the intensity measurements are interpolated from the measured m/z array to the shared m/z array.

4 RESULTS AND DISCUSSION

4.1 Intrasample mass alignment

In this section, three methods are applied to the MALDI-MSI dataset described in Section 3: (i) the proposed method without optimization; (ii) the proposed method with optimization; (iii) the publicly available MSIWarp (Eriksson et al., 2020). The samples of the dataset are labeled as A-X. To evaluate the performance of the proposed method, the aforementioned methods are applied to the dataset and compared based on: mass dispersion, cosine similarity and timings. The mean mass dispersion and cosine similarity per sample of the dataset before and after alignment is shown in Figure 3, in conjunction with relevant timing information. The results are summa-

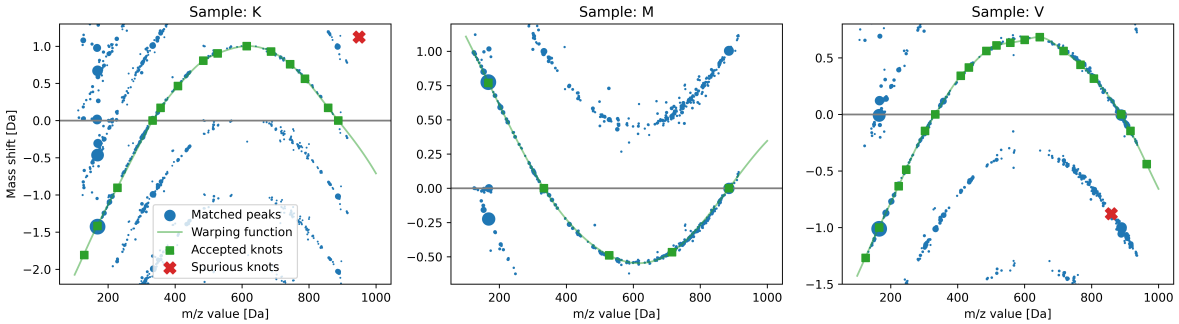

Figure 4: Initial estimates of the warping function for mass spectra from sample K, M and V which have large nonlinear mass shifts compared to the reference mass spectrum. Spurious knots are detected and removed by the outlier detection procedure.

rized in Table 1. Table S2 contains the parameters used for the proposed method. Table S3 summarizes the timings of the experiments for the different algorithms. The cosine similarity, the mean and median mass dispersion of the 100 most intense peaks of the average spectrum are reported in Table S4. The results before solving the optimization problem are also reported. The additional processing and parameter selection for MSIWarp are discussed in the Supporting Information A.1.

In most cases, the proposed algorithm with and without optimization significantly reduces the mass dispersion of the samples. Optimization further improves the initial estimates of the knots by maximizing the cosine similarity, which is associated with a reduction of the median mass dispersion, except for samples A and Q. The mean mass dispersion and the cosine similarity after alignment have improved for all samples (Table S4). If the local warping estimates are chosen correctly, then they will have a maximum error of half of the local m/z bin width. Solving the subsequent optimization problem can drastically lower the error and improve the results if the initial estimate is in the neighborhood of the optimal warping function. The larger the error, the higher the probability of converging to a suboptimal warping function. Evidently, the number of knots must be sufficiently large to model the complexity of the unknown warping function.

Table 1: Summary of the results reported in Figure 3. The mass dispersion reported is the median of the median mass dispersion across all samples. The cosine similarity reported is the median of the mean cosine similarities across all samples.

Metric	Mass dispersion	Cosine
Original	26.62 ppm	0.9021
No optimization	16.34 ppm	0.9368
Optimization	12.41 ppm	0.9524
MSIWarp	12.84 ppm	0.9466

The performance of the proposed method with optimization is slightly better than the performance of MSIWarp, whereas our method without optimization is outperformed by MSIWarp and our algorithm with optimization. Table 1 shows that the median mass dispersion across the dataset is reduced from 26.62 ppm to 16.34 ppm before optimization, and it is reduced to 12.41 ppm after optimization and to 12.84 ppm with MSIWarp. The median cosine similarity is improved from 0.9021 to 0.9368 before optimization, 0.9524 after optimization, and to 0.9466 after applying MSIWarp. Since the cosine similarity is optimized for every mass spectrum in the sample, it is not surprising that the proposed method with optimization has the most significant increase in this metric. Although the performance metrics of our method without optimization are worse than the method with optimization and MSIWarp, the execution time is found to be significantly faster than these methods, while still achieving a considerable reduction in mass misalignment. These characteristics are useful for fast alignment of typically large MSI datasets.

The performance metrics for samples K, M and V are improved substantially for both the method with and without optimization, whereas MSIWarp is unable to reduce the mass misalignment for these samples. Figure 4 displays the warping of one spectrum per sample to the reference spectrum using the proposed method without optimization. Evidently, the mass shifts for some pixels in these samples are severe and our algorithm is able to correct these large errors consistently, highlighting the robustness of the developed approach.

It is already established that the proposed method without optimization is considerably faster than the other approaches considered. Furthermore, experiments suggest that the proposed method with optimization is roughly three to four times as fast as MSIWarp. The computational complexity of the method is linear with respect to the number of mass spectra

present in the sample and the number of segments considered (Byrd et al., 1995). As the number of knots increases, they influence a smaller portion of the m/z values. This can be incorporated into the evaluation of the gradient to speed up the algorithm.

4.2 Parameter selection

The proposed peak-shape aware approach is characterized by several parameters: (i) the number of segments into which the mass spectra are segmented; (ii) the size of the extracted window around the maximum peak of each segment; (iii) the maximal mass shift allowed while computing the estimate expressed as a factor which is multiplied with the extracted window size; (iv) whether to perform the outlier detection procedure, and if so, its residual threshold; and finally (v) whether to solve the optimization problem, and if so, the maximum allowed mass shift Δ. Understanding the purpose of these parameters is important in order to obtain desirable results: parameter (i) influences the complexity of the warping function; (ii) and (iii) determine the maximal mass shift considered; (iv) prevents incorrect alignment; and (v) optimizes the knots.

The effects of the parameters on alignment are assessed by performing several experiments with various parameters on samples K and L. These samples were chosen based on their limited size and the severity of mass shifts observed: sample K exhibits severe mass shifts, while mass shifts of L are mild. The results of the experiments are summarized in Tables S5 and S6.

Usually, the expected amount of mass shift is known and parameters (ii) and (iii) can be chosen accordingly, eliminating one degree of freedom. In the experiment, the parameters are chosen such that the mass shift considered for the determination of the local alignment is identical for sample K or L. This allows for the assessment of the impact of varying (ii): the window size around the maximum peak per segment. If severe mass shifts are present, the correct determination of the local mass shift is expected to be harder when the window size increases, since it is more likely that the local mass shift is not appropriate for peaks in proximity to the window boundaries. This expectation is confirmed by the results of the experiment for sample K with severe mass shifts. On the contrary, for sample L with mild mass shifts this effect is negligible.

Experiments were also conducted with various numbers of segments in which the mass spectra are divided. Eventually, this results in the number of knots and the size of the optimization problem after an optional outlier detection step. Performance metrics usually improve as the number of segments increases, because an increase in knots means that the spline function can approximate more complex warping functions. Additionally, the segments become smaller and may only contain peaks that are not present in both the mass spectrum and the reference spectrum, resulting in more outlier knots. In the event that multiple outlier knots exist in close proximity, true knots might be removed. However, outlier detection proves to be beneficial if the number of segments is not exceedingly small or large. For example, applying outlier detection to a sample with severe mass shifts and five segments deteriorates the performance metrics since the local mass shifts are far from linear, see Figure 4. Nevertheless, applying the algorithm with outlier detection to samples that exhibit minor mass shifts is inconsequential since erroneous matches are already eliminated by a good choice of (ii) and (iii). Furthermore, the execution time hardly increases if the number of segments increases and the optimization problem is not solved. Conversely, the execution time increases sublinearly with the number of knots if the optimization problem is solved, thanks to the efficient computation of the gradient mentioned in the previous section.

Now that we know the effect of the parameters is explained, we propose the following strategy to perform parameter selection. The performance of our algorithm without optimization is a strong indicator of the performance of the algorithm with optimization, since the performance metrics for the former are empirically observed to be a lower bound for the latter, see Figure 3. Because of this, the proposed method without optimization can be used to perform parameter optimization on a subset of the data, significantly speeding up the process.

4.3 Multimodal mass alignment

Aside from solving mass misalignment within a sample of a single modality, the proposed method can also align mass spectra from different MS modalities, provided that there is a sufficient amount of shared information between the modalities. In this section, our method is applied to align the corrected MALDI-MSI and REIMS measurements by computing a single warping function from the average MALDI-MSI mass spectrum to the average REIMS mass spectrum per sample. The developed method expects the m/z array of the moving spectrum and the reference spectrum to be identical. Consequently, the REIMS measurements are binned to the m/z array of the MALDI-TOF measurements since the REIMS

data has a higher spectral resolution.

As the lower part of the m/z range is dominated by matrix-related peaks that are not shared between the modalities, cross-modal alignment is performed on the 600–900 Da m/z range. The distribution of the resulting cosine similarity before and after alignment is shown in Figure 5. The average similarity improves from 0.101 to 0.180 after alignment using the initial estimates as knots for the warping function, and to 0.189 after optimizing the knots. The cosine similarity between the mass spectra of MALDI-MSI and REIMS roughly doubles after alignment. The relatively low cosine similarity values suggest that the amount of shared information is limited. Abundant peaks in one modality may not be present in the other, and the shapes of the peaks may differ between the modalities. Notwithstanding, the proposed methodology can considerably improve the metric.

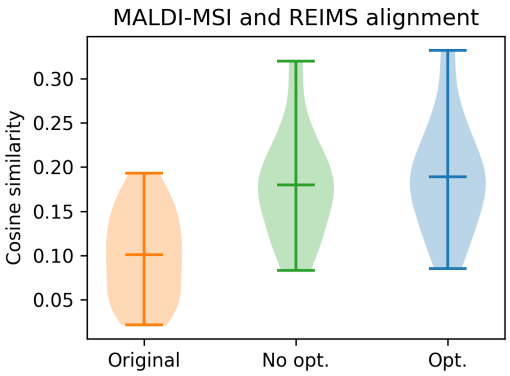

Figure 5: Violinplot of the cosine similarity before and after alignment between the average MALDI-MSI mass spectrum and the average REIMS mass spectrum of a sample in the 600–900 Da range.

4.4 Intersample mass alignment

Naturally, it is also possible to align the measurements across the samples. The average REIMS measurement across the tissues is used as a reference spectrum. After alignment of the REIMS measurements, the median mass dispersion of the 50 largest peaks improves from 12.15 ppm to 6.70 ppm without optimization and to 6.08 ppm with optimization. Similarly, the MALDI-MSI misalignment between samples can be reduced by computing a warping function from the average mass spectra of the samples to a reference spectrum which can be chosen as the average mass spectrum of any sample. Then, the median mass dispersion of the 100 largest peaks of these average mass spectra improves from 105.10 ppm to 23.60 ppm without optimization and to 19.48 ppm with optimization. Finally, the resulting average MALDI-MSI mass spectra can be aligned with the resulting average REIMS mass spectra. This indirect approach to perform multimodal mass alignment improves the average similarity from 0.054 to 0.178 without optimization and to 0.181 with optimization.

4.5 General applicability to spectral data

The main drawback of the proposed alignment approach with respect to applicability is the requirement of profile data. Consequently, the method is not applicable to centroided datasets. In the case of sparse mass spectrometry data, the algorithm can be applied after careful construction of a shared m/z array, as described in Section 3.

Furthermore, the proposed methodology is applicable to any modality that can benefit from alignment assuming that an initial estimate close to the optimal warping function can be determined reliably. To illustrate this general applicability, the method can be applied to spectral data with a limited number of peaks. For example, the proposed method is able to align 1D Nuclear Magnetic Resonance (NMR) metabolome profiles of table wines (Larsen et al., 2006). CluPA (Vu et al., 2011) is a method introduced specifically to align NMR data and improves the mean cosine similarity of the spectra with the average spectrum of 0.8486 to 0.9945, whereas the proposed algorithm without optimization improves the metric to 0.9946 and with optimization to 0.9968. This shows the general applicability of the proposed algorithm.

4.6 Implementation details

The described algorithm is implemented in Python using Numba (Lam et al., 2015) to improve its execution speed. The optimization algorithm L-BFGS-B is implemented using SciPy (Virtanen et al., 2020). Since the proposed algorithm warps every pixel independently, this is done in parallel. The experiments were performed on a laptop with an AMD Ryzen 7 PRO 7840HS CPU (3.8 GHz) that contains 8 cores and 32 GB RAM. The source code is available on GitHub: https://github.com/VanhemelThomas/psalign.

5 CONCLUSIONS

This work proposes a fast and robust data-driven algorithm to reduce the mass misalignment between mass spectra. First, a coarse alignment is com-

puted based on the profile data and subsequently refined by solving a mathematical optimization problem to maximize the similarity between the mass spectra. The procedure does not require peak-picking: the subtle features of profile spectra, such as the peak shape, contribute to the robust determination of the warping function. The proposed methodology is evaluated by performing extensive experiments on MALDI-TOF MSI and REIMS datasets and the performance metrics indicate that the refined alignment performs slightly better than MSIWarp, while the efficient coarse and refined alignment procedures reliably correct extreme mass shifts. The developed approach can perform intrasample, intersample and multimodal alignment, further demonstrating its robustness. Furthermore, the proposed methodology is shown to be generally applicable to any spectral data modality that exhibits sufficiently distinct spectral regions. By ensuring consistent mass alignment across mass spectra, our method enhances the integrity of downstream data analysis, which is especially important when working directly with profile data that retains subtle but informative spectral features.

ACKNOWLEDGEMENTS

This work was supported in part by the KU Leuven: Research Fund (IDN/24/013, iBOF/23/064), Industrial Research Fund and several Leuven Research and Development bilateral industrial projects; in part by Flemish Government agencies: FWO (TBM T001725N, TBM T001919N, PhD Grant 11K5623N), EWI, VLAIO (CSBO HBC.2021.0076); in part by the FOD VVVL (AI-943); in part by Interreg Vlaanderen-Nederland (Molecular Brain Tumor Detector project); in part by the Dutch Province of Limburg through Interreg Vl-Nl and the LINK program; in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 885682). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or ERC. Neither the European Union nor the granting authority can be held responsible for them. The authors would like to gratefully acknowledge the help from the neurosurgeons at UZ Leuven, Belgium for providing tissue samples. The resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government. The authors declare no competing financial interest.

REFERENCES

Abdelmoula, W. M., Lopez, B. G.-C., Randall, E. C., Kapur, T., Sarkaria, J. N., White, F. M., Agar, J. N., Wells, W. M., and Agar, N. Y. R. (2021). Peak learning of mass spectrometry imaging data using artificial neural networks. *Nature Communications*, 12(1):5544.

Balog, J., Szaniszlo, T., Schaefer, K.-C., Denes, J., Lopata, A., Godorhazy, L., Szalay, D., Balogh, L., Sasi-Szabo, L., Toth, M., and Takats, Z. (2010). Identification of Biological Tissues by Rapid Evaporative Ionization Mass Spectrometry. *Analytical Chemistry*, 82(17):7343–7350.

Bocker, S. and Makinen, V. (2008). Combinatorial Approaches for Mass Spectra Recalibration. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 5(1):91–100.

Boor, C. D. (1978). *A Practical Guide to Splines*. Springer New York.

Boskamp, T., Lachmund, D., Casadonte, R., Hauberg-Lotte, L., Kobarg, J. H., Kriegsmann, J., and Maass, P. (2020). Using the Chemical Noise Background in MALDI Mass Spectrometry Imaging for Mass Alignment and Calibration. *Analytical Chemistry*, 92(1):1301–1308.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Caprioli, R. M., Farmer, T. B., and Gile, J. (1997). Molecular Imaging of Biological Samples: Localization of Peptides and Proteins Using MALDI-TOF MS. *Analytical Chemistry*, 69(23):4751–4760.

Denker, A., Behrmann, J., and Boskamp, T. (2024). Improved Mass Calibration in MALDI MSI Using Neural Network-Based Recalibration. *Analytical Chemistry*, 96(19):7542–7549.

Eriksson, J. O., Sánchez Brotons, A., Rezeli, M., Suits, F., Markó-Varga, G., and Horvatovich, P. (2020). MSIWarp: A General Approach to Mass Alignment in Mass Spectrometry Imaging. *Analytical Chemistry*, 92(24):16138–16148.

Fischler, M. A. and Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395.

Gobom, J., Mueller, M., Egelhofer, V., Theiss, D., Lehrach, H., and Nordhoff, E. (2002). A Calibration Method That Simplifies and Improves Accurate Determination of Peptide Molecular Masses by MALDI-TOF MS. *Analytical Chemistry*, 74(15):3915–3923.

Gross, J. H. (2017). *Mass Spectrometry*. Springer International Publishing, Cham.

Hendriks, T. F., Birmpili, A., de Vleeschouwer, S., Heeren, R. M., and Cuypers, E. (2025). Integrating Rapid Evaporative Ionization Mass Spectrometry Classification with Matrix-Assisted Laser Desorption Ionization Mass Spectrometry Imaging and Liquid Chromatography-Tandem Mass Spectrometry to

Unveil Glioblastoma Overall Survival Prediction. *ACS Chemical Neuroscience*, 16(6):1021–1033.

Kulkarni, P., Kaftan, F., Kynast, P., Svatoš, A., and Böcker, S. (2015). Correcting mass shifts: A lock mass-free recalibration procedure for mass spectrometry imaging data. *Analytical and Bioanalytical Chemistry*, 407(25):7603–7613.

La Rocca, R., Kune, C., Tiquet, M., Stuart, L., Eppe, G., Alexandrov, T., De Pauw, E., and Quinton, L. (2021). Adaptive Pixel Mass Recalibration for Mass Spectrometry Imaging Based on Locally Endogenous Biological Signals. *Analytical Chemistry*, 93(8):4066–4074.

Lam, S. K., Pitrou, A., and Seibert, S. (2015). Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, LLVM '15, pages 1–6, New York, NY, USA. Association for Computing Machinery.

Larsen, F. H., van den Berg, F., and Engelsen, S. B. (2006). An exploratory chemometric study of 1h nmr spectra of table wines. *Journal of Chemometrics*, 20(5):198–208.

McCann, A., Rappe, S., La Rocca, R., Tiquet, M., Quinton, L., Eppe, G., Far, J., De Pauw, E., and Kune, C. (2021). Mass shift in mass spectrometry imaging: Comprehensive analysis and practical corrective workflow. *Analytical and Bioanalytical Chemistry*, 413(10):2831–2844.

Norris, J. L., Cornett, D. S., Mobley, J. A., Andersson, M., Seeley, E. H., Chaurand, P., and Caprioli, R. M. (2007). Processing MALDI Mass Spectra to Improve Mass Spectral Direct Tissue Analysis. *International journal of mass spectrometry*, 260(2-3):212–221.

Ràfols, P., Castillo, E. D., Yanes, O., Brezmes, J., and Correig, X. (2018). Novel automated workflow for spectral alignment and mass calibration in MS imaging using a sputtered Ag nanolayer. *Analytica Chimica Acta*, 1022:61–69.

Robichaud, G., Garrard, K. P., Barry, J. A., and Muddiman, D. C. (2013). Msireader: An open-source interface to view and analyze high resolving power ms imaging files on matlab platform. *Journal of the American Society for Mass Spectrometry*, 24(5):718–721. PMID: 23536269.

Smets, T., Verbeeck, N., Claesen, M., Asperger, A., Griffioen, G., Tousseyn, T., Waelput, W., Waelkens, E., and De Moor, B. (2019). Evaluation of Distance Metrics and Spatial Autocorrelation in Uniform Manifold Approximation and Projection Applied to Mass Spectrometry Imaging Data. *Analytical Chemistry*, 91(9):5706–5714.

Tracy, M. B., Chen, H., Weaver, D. M., Malyarenko, D. I., Sasinowski, M., Cazares, L. H., Drake, R. R., Semmes, O. J., Tracy, E. R., and Cooke, W. E. (2008). Precision enhancement of MALDI-TOF MS using high resolution peak detection and label-free alignment. *Proteomics*, 8(8):1530–1538.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.

Vu, T. N., Valkenborg, D., Smets, K., Verwaest, K. A., Dommisse, R., Lemière, F., Verschoren, A., Goethals, B., and Laukens, K. (2011). An integrated workflow for robust alignment and simplified quantitative analysis of NMR spectrometry data. *BMC Bioinformatics*, 12:405.

Wijetunge, C. D., Saeed, I., Boughton, B. A., Roessner, U., and Halgamuge, S. K. (2015). A new peak detection algorithm for MALDI mass spectrometry data based on a modified Asymmetric Pseudo-Voigt model. *BMC Genomics*, 16(12):S12.

Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. (1997). Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560.

A.1 MSIWarp preprocessing and parameters

The mass spectra were peak-picked, after signal smoothing with the Savitzky-Golay filter, using `find_peaks` from the SciPy package with the `distance` and `height` parameters chosen to be 30 and the 90th percentile of intensities, respectively, which result in on average 378 detected peaks per mass spectrum per sample. Exceptions are R, H and Q with a `distance` value of 20, and D and L with a `height` value of the 80th percentile of intensities. The parabolic centroiding algorithm was used to obtain an accurate estimate of the peak mass. Experiments were performed with different reference spectra choices and the results of the best alignment per sample were used in the generation of Figure 3. The reference spectra considered were the mass spectrum with maximal TIC and the spectrum at index 200. If the spectrum at index 200 did not contain sufficient shared information, it was replaced with index 2000, 2500 or 21347. The parameters for the uniform node placement were `n_steps` $= 50$, `n_peaks` $= 30$ and `max_nodes` $= 8$. σ was chosen as 150 ppm and ε as 765 ppm for all samples except K, M and V due to the larger mass shifts present in these samples, for these samples additional experiments with $\varepsilon = 15300$ ppm and RANSAC, with parameters `n_segments` $= 2$, `n_iterations` $= 50$, `n_samples` $= 2$, `min_n_pairs` $= 10$, `n_steps` $= 25$ and `distance_threshold` $= 1$, were performed. Finally, the computed warping functions are applied to the profile data and the performance metrics are calculated based on the results.

Table S1: Detailed information of the MALDI-TOF samples used in this paper.

Sample	A	B	C	D	E	F	G	H
Number of m/z bins	36866	36866	36854	36856	36849	36846	36866	27397
Number of pixels	28114	11879	37423	16784	21497	31208	27036	62521
Max bin width [Da]	0.024	0.024	0.024	0.024	0.024	0.024	0.024	0.033

Sample	I	J	K	L	M	N	O	P
Number of m/z bins	36868	36868	36841	36866	38581	36049	36045	36049
Number of pixels	83446	62931	23122	12559	42778	42400	78953	56924
Max bin width [Da]	0.024	0.024	0.024	0.024	0.023	0.022	0.025	0.022

Sample	Q	R	S	T	U	V	W	X
Number of m/z bins	27398	27398	36868	36867	36868	36908	36868	38597
Number of pixels	110620	28831	25852	80885	59834	46576	18330	77010
Max bin width [Da]	0.033	0.033	0.024	0.024	0.024	0.024	0.024	0.023

Table S2: Parameters used by the proposed algorithm to generate Figure 3. The residual threshold used in the outlier detection procedure is the mean absolute deviation clipped at 0.1 and 0.3. The reference spectra used were the average mass spectra after TIC normalization, with the exception of samples K, M and V for which the mass spectra of pixel 200, 2500 and 21347 respectively were employed.

Sample	A	B	C	D	E	F	G	H
Segments	5	5	10	5	5	15	3	16
Window	2000	2000	500	2000	2000	400	2000	500
Factor	1.03	1.03	1.08	1.03	1.03	1.05	1.03	1.02
Outlier removal	Yes	No	Yes	No	No	Yes	No	Yes

Sample	I	J	K	L	M	N	O	P
Segments	5	5	15	5	5	5	5	5
Window	2000	1000	200	500	500	2000	2000	500
Factor	1.03	1.005	1.4	1.12	1.12	1.01	1.03	1.02
Outlier removal	Yes	No	Yes	Yes	No	No	Yes	No

Sample	Q	R	S	T	U	V	W	X
Segments	3	3	5	3	5	20	5	5
Window	2000	2000	2000	500	2000	100	2000	2000
Factor	1.01	1.01	1.03	1.12	1.02	2	1.02	1.03
Outlier removal	No	No	Yes	No	Yes	Yes	Yes	Yes

Table S3: Execution time of the algorithms applied to samples of the MALDI-TOF dataset. These results were used to generate Figure 3.

Sample		A	B	C	D	E	F	G	H
Time [s]	No optimization	16.10	6.72	19.16	9.23	11.96	18.22	14.27	34.75
	Optimization	191	83	315	108	141	346	159	469
	MSIWarp	592	294	1146	682	522	768	713	1334

Sample		I	J	K	L	M	N	O	P
Time [s]	No optimization	58.40	32.54	12.40	7.19	22.40	22.12	43.92	27.91
	Optimization	542	393	295	78	350	278	584	392
	MSIWarp	2253	1386	1215	424	1276	1144	2208	1498

Sample		Q	R	S	T	U	V	W	X
Time [s]	No optimization	50.73	12.79	14.55	52.98	31.89	26.19	9.95	48.16
	Optimization	496	123	161	440	363	591	124	588
	MSIWarp	2507	708	595	2451	1639	1830	435	2065

Table S4: Numerical results after applying the algorithms to the MALDI-TOF dataset. The mean and median mass dispersion and the cosine similarity are reported. These results were used to generate Figure 3.

Sample		A	B	C	D	E	F	G	H
Mean mass disp. [ppm]	Original	12.67	16.25	14.51	20.54	19.36	27.36	23.53	27.60
	No opt.	12.00	15.39	14.52	18.99	19.60	22.94	17.79	21.48
	Opt.	11.19	12.26	11.54	13.91	17.06	15.98	10.92	20.12
	MSIWarp	13.56	11.77	15.26	18.84	19.36	21.64	15.09	18.63
Median mass disp. [ppm]	Original	10.65	12.04	13.01	15.57	15.92	16.35	21.45	22.12
	No opt.	9.66	10.50	12.50	14.84	15.84	16.00	14.40	15.64
	Opt.	9.87	7.47	9.45	11.99	14.49	12.19	8.45	13.60
	MSIWarp	12.48	8.39	10.19	13.71	15.92	12.43	9.39	13.97
Cosine similarity	Original	0.957	0.960	0.973	0.925	0.925	0.947	0.930	0.934
	No opt.	0.958	0.962	0.973	0.928	0.927	0.957	0.942	0.947
	Opt.	0.965	0.973	0.977	0.937	0.936	0.969	0.956	0.969
	MSIWarp	0.961	0.973	0.976	0.931	0.925	0.962	0.952	0.973

Sample		I	J	K	L	M	N	O	P
Mean mass disp. [ppm]	Original	24.62	26.38	35.56	29.51	36.27	34.20	39.99	38.87
	No opt.	17.37	25.04	20.83	16.88	18.94	23.20	23.19	26.19
	Opt.	12.46	21.05	13.70	11.01	14.67	16.88	17.81	18.31
	MSIWarp	20.46	21.49	35.88	13.70	34.82	16.39	29.33	14.47
Median mass disp. [ppm]	Original	23.71	23.75	24.69	25.91	27.34	28.57	32.59	32.72
	No opt.	14.87	20.76	15.51	13.14	14.71	17.42	17.88	18.24
	Opt.	11.01	17.23	11.31	8.21	12.03	13.40	15.96	13.31
	MSIWarp	13.19	17.79	23.45	7.99	25.39	11.59	17.59	9.01
Cosine similarity	Original	0.929	0.964	0.882	0.905	0.812	0.899	0.846	0.876
	No opt.	0.939	0.968	0.969	0.938	0.945	0.926	0.888	0.914
	Opt.	0.952	0.981	0.978	0.952	0.952	0.952	0.917	0.957
	MSIWarp	0.941	0.977	0.889	0.949	0.820	0.950	0.862	0.953

Sample		Q	R	S	T	U	V	W	X
Mean mass disp. [ppm]	Original	39.23	50.13	41.70	52.88	54.33	66.97	95.92	108.25
	No opt.	27.98	23.78	28.08	23.26	27.08	24.07	28.19	26.66
	Opt.	25.30	16.14	17.51	12.29	17.74	13.86	17.39	17.49
	MSIWarp	26.07	22.40	19.28	14.75	17.99	66.97	22.00	23.05
Median mass disp. [ppm]	Original	34.59	38.32	38.87	49.41	49.92	50.14	87.38	103.05
	No opt.	21.70	16.69	21.16	17.26	21.96	18.62	23.65	20.25
	Opt.	23.82	12.63	14.63	10.23	15.46	10.97	14.02	15.02
	MSIWarp	19.85	11.04	11.90	9.30	11.51	50.14	14.87	15.84
Cosine similarity	Original	0.895	0.870	0.908	0.893	0.877	0.635	0.855	0.830
	No opt.	0.912	0.935	0.930	0.937	0.927	0.934	0.940	0.937
	Opt.	0.949	0.956	0.953	0.951	0.947	0.949	0.954	0.948
	MSIWarp	0.933	0.945	0.945	0.948	0.943	0.635	0.949	0.940

Table S5: Median mass dispersion and execution time of applying the proposed algorithm with and without solving the optimization problem for various parameters to sample K. The first two row represent the window and the factor parameter respectively, and the first column represents the number of segments. Outlier detection was used for the rows which contain a star (*) next to the number of segments. The residual threshold used in the outlier detection procedure is the mean absolute deviation clipped at 0.1 and 0.3. Entries with value - violate the constraint: factor $\times$ window $\geq \frac{\text{Number of m/z bins}}{2\times\text{segments}}$.

(a) No optimization: median mass dispersion in ppm.

Window	200	500	1000	2000
Factor	1.4	1.16	1.08	1.04
Segments				
5	16.45	16.57	16.27	16.09
5*	23.98	23.95	24.19	24.05
10*	16.13	18.72	19.51	-
15*	15.51	18.29	22.14	-
20*	15.98	20.04	-	-

(b) Optimization: median mass dispersion in ppm.

Window	200	500	1000	2000
Factor	1.4	1.16	1.08	1.04
Segments				
5	14.01	14.52	13.92	13.11
5*	21.31	21.31	23.30	23.75
10*	12.46	14.78	18.06	-
15*	11.31	11.40	20.25	-
20*	11.63	17.24	-	-

(c) No optimization: time in *s*.

Window	200	500	1000	2000
Factor	1.4	1.16	1.08	1.04
Segments				
5	14.58	13.28	13.16	12.68
5*	13.13	12.60	13.22	13.22
10*	12.54	12.75	15.18	-
15*	13.08	13.56	15.87	-
20*	13.37	14.60	-	-

(d) Optimization: time in *s*.

Window	200	500	1000	2000
Factor	1.4	1.16	1.08	1.04
Segments				
5	155	147	146	148
5*	134	134	131	132
10*	194	206	198	-
15*	251	260	236	-
20*	279	283	-	-

Table S6: Median mass dispersion and execution time of applying the proposed algorithm with and without solving the optimization problem for various parameters to sample L. The first two row represent the window and the factor parameter respectively, and the first column represents the number of segments. Outlier detection was used for the rows which contain a star (*) next to the number of segments. The residual threshold used in the outlier detection procedure is the mean absolute deviation clipped at 0.1 and 0.3. Entries with value - violate the constraint: factor $\times$ window $\geq \frac{\text{Number of m/z bins}}{2\times\text{segments}}$.

(a) No optimization: median mass dispersion in ppm.

Window	200	500	1000	2000
Factor	1.3	1.12	1.06	1.03
Segments				
3	12.32	12.14	12.10	12.03
5	13.53	13.15	12.78	12.49
5*	13.52	13.15	12.78	12.49
10*	13.15	12.53	12.24	-
20*	13.78	12.70	-	-

(b) Optimization: median mass dispersion in ppm.

Window	200	500	1000	2000
Factor	1.3	1.12	1.06	1.03
Segments				
3	8.13	8.14	8.20	8.21
5	8.30	8.21	8.23	8.52
5*	8.30	8.20	8.22	8.52
10*	8.17	8.12	8.22	-
20*	9.12	8.58	-	-

(c) No optimization: time in *s*.

Window	200	500	1000	2000
Factor	1.3	1.12	1.06	1.03
Segments				
3	7.08	6.75	6.67	7.26
5	8.72	7.16	7.50	7.89
5*	7.46	7.99	7.13	8.40
10*	7.85	7.88	8.80	-
20*	8.24	8.72	-	-

(d) Optimization: time in *s*.

Window	200	500	1000	2000
Factor	1.3	1.12	1.06	1.03
Segments				
3	73	72	74	73
5	89	89	88	88
5*	91	89	89	88
10*	115	110	112	-
20*	162	160	-	-