Analytica Chimica Acta 1230 (2022) 340247

Contents lists available at ScienceDirect

Analytica Chimica Acta

journal homepage: www.elsevier.com/locate/aca

Incorporating measurement variability when comparing sets of high-resolution mass spectra

Matthew J. Roberts, Arun S. Moorthy, Edward Sisco, Anthony J. Kearsley *

National Institute of Standards and Technology, Gaithersburg, MD, 20899, USA

HIGHLIGHTS

- High dimensional consensus (HDC) mass spectra: a novel formulation for describing mass spectra as probability distributions.
- Mathematical framework for comparing HDC mass spectra: after embedding mass spectra into an infinite dimensional Hilbert space, establish affordable analog of commonly employed similarity measure.
- Performance: structurally similar compounds not differentiable with standard similarity are distinguishable by HDC if measurements collected under similar conditions.

GRAPHICAL ABSTRACT

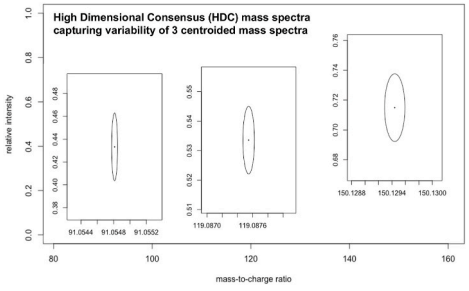

ARTICLE INFO

Keywords:
Sample discrimination
High-dimensional consensus (HDC) mass spectra
High-resolution mass spectrometry
Mass spectral similarity
Mass spectrometry
Probability distributions

ABSTRACT

Mass spectra are an important signature by which compounds can be identified. We recently formulated a mathematical approach for incorporating measurement variability when comparing sets of high-resolution mass spectra. Leveraging replicate mass spectra, we construct high-dimensional consensus mass spectra—representing each of the compared analytes—and compute the similarity between these data structures. In this paper, we present this approach and discuss its applications and limitations when trying to discriminate methamphetamine and phentermine using in-source collision induced dissociation mass spectra collected with direct analysis in real time mass spectrometry.

1. Introduction

Comparing mass spectra of unidentified compounds to reference mass spectra of known compounds is a key step in the compound identification process [1,2]. Based on this comparison—which is done both visually and numerically using similarity scores—*an analyst must decide* whether the mass spectra are sufficiently similar (or dissimilar) to confirm (or deny) that the spectra are measurements of the same compounds [3].

Fig. 1 shows the chemical structures of methamphetamine (top/black) and phentermine (bottom/blue), as well as head-to-tail plots of representative high-resolution *in-source collision induced dissociation* (is-CID) mass spectra collected with *Direct Analysis in Real Time Mass Spectrometry* (DART-MS). The resulting mass spectra are more likely to present stable protonated molecules with lower fragmentation voltages (e.g., +30 V orifice 1 voltage) and are increasingly likely to show

* Corresponding author.
E-mail address: Anthony.Kearsley@nist.gov (A.J. Kearsley).

https://doi.org/10.1016/j.aca.2022.340247
Received 16 May 2022; Received in revised form 4 August 2022; Accepted 7 August 2022
Available online 15 August 2022

fragment ions with increasing fragmentation voltages (e.g. +60 and +90 V orifice 1 voltage). In this figure, we see that the +30 V is-CID mass spectra of methamphetamine and phentermine are visually dissimilar, and a standard cosine similarity between these mass spectra is 0.683. We also see that the +60 V and +90 V spectra of methamphetamine and phentermine are both dominated by a single fragment ion (nominal *m/z* 91 corresponding to a tropylium cation), and are visually *difficult* and numerically impossible to distinguish. There are minor ions that may help an analyst visually conclude that the spectra are measurements of two different compounds—especially if the analyst is approaching the spectra with the intention of finding differences—however, with mass spectra collected through DART-MS or other ambient ionization MS techniques, it is possible that these minor peaks are non-reproducible artifacts.

Although an analyst is unlikely to be trying to discriminate methamphetamine and phentermine using only *high fragmentation* is-CID mass spectra collected with an ambient ionization MS technique, the problem is exemplary of a more general challenge across mass spectral interpretation; many unique compounds produce near-identical mass spectra that are difficult to *objectively discriminate* with current approaches.

Cosine similarity is one of several numerical *similarity scores*, or *match factors*, that have been discussed in the literature as objective metrics by which an analyst can discriminate mass spectra [4–14]. Most of these similarity scores were designed (i) for low resolution mass spectra where mass-to-charge ratios are mapped to nominal mass, and (ii) in the context of mass spectral library searching, where a single query mass spectrum is compared sequentially to individual reference mass spectra. However, we know that replicate high-resolution mass spectra of analytes will often have observable variability in both mass-to-charge ratios and relative intensities [15,16], and so any traditional similarity score computed with a pair of single high-resolution measurements will implicitly include some level of uncertainty.

We recently formulated a new mathematical approach for comparing two samples when provided with replicate high-resolution mass spectra. Using these replicate measurements, we are able to consider peak variability when computing a *high-dimensional consensus (HDC) mass spectral similarity score*. The approach begins with constructing an HDC mass spectrum data structure with the available replicate measurements, and follows with comparing two of these data structures, each representing one of the samples. While this new approach, by virtue of design, must provide a more *complete characterization* of the similarity between sets of high-resolution mass spectra, a natural question is *can we use peak variability differences as a signature to discriminate near identical mass spectra?* such as the +60 V is-CID mass spectra shown in Fig. 1. In this paper, we present our first formal description of computing HDC similarity scores between high-resolution mass spectra and report our preliminary observations from applying the method to the example problem in Fig. 1.

2. Methods

2.1. Computing high-dimensional consensus (HDC) mass spectral similarity scores

This method of spectral comparison can best be described in two distinct steps: (i) building a data structure that captures the variability observed in a set of replicate high-resolution mass spectra, and (ii) computing a measure of similarity between a pair of these data structures.

The term *consensus mass spectrum* is commonly used to refer to a data structure constructed from "averaging" mass spectra collected over several scans in the same experiment [17]. There are a variety of ways to

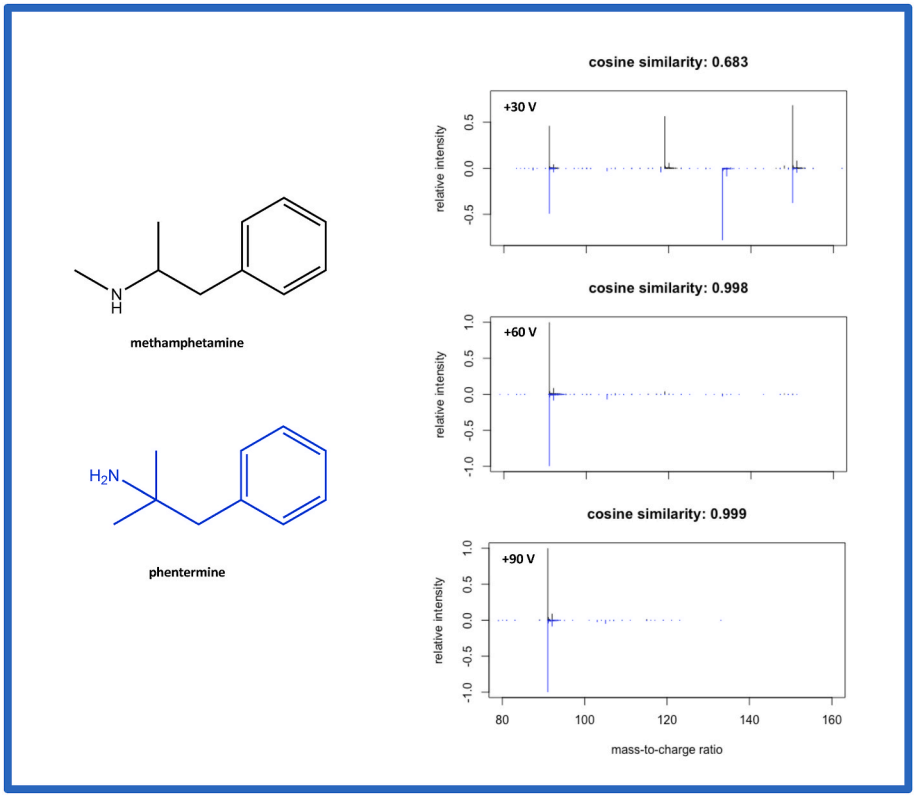

Fig. 1. Chemical structures (left) and representative in-source collision induced dissociation (is-CID) mass spectra of methamphetamine and phentermine (right). The is-CID mass spectra were collected with DART-MS using orifice 1 voltages set to +30 V, +60 V, and +90 V. In the head-to-tail plots of mass spectra, the top/black mass spectra are measurements of methamphetamine and the bottom/blue mass spectra are measurements of phentermine.

average mass spectra, but the result is always a set of coupled *average* mass-to-charge ratios (x) and *average* relative intensity values (y); the structure of the constructed consensus mass spectrum is thus identical to any measured mass spectrum and so consensus mass spectra can be readily inserted into traditional algorithms for estimating mass spectral similarity.

In our new approach, we construct *high-dimensional consensus* (HDC) mass spectra, which is a consensus mass spectrum together with the associated standard deviation of the mass-to-charge ratios and relative peak intensities (i.e., a set of points in $\mathbb{R}^4$ as opposed to $\mathbb{R}^2$). These high-dimensional data structures are a natural extension of traditional consensus mass spectra, carrying additional information about spectral variability that is typically discarded in a consensus mass spectrum.

The process of constructing an HDC mass spectrum can be summarized as follows: We begin with a collection of n replicate mass spectra. We normalize the relative intensities in each mass spectrum by dividing each intensity value by the square root of the sum of squares of all intensity values in the spectrum (L2-norm); given that the method relies on observed variability, it is important that we do not normalize to a base peak, thus artificially removing variability from the most abundant peak across the spectra. Then, beginning with the most intense peak of all the measurements, p_1, we identify the single peak in each replicate measurement that is closest to p_1 when considering both m/z and relative intensity values (euclidean distance). This set of n peaks, which we can refer to as S_1, is used to define the first peak statistic $P_1 = (\overline{p}_{x,1}, \overline{p}_{y,1}, s_{p_{x,1}}, s_{p_{y,1}})$, where $\overline{p}_{x,1}$ and $\overline{p}_{y,1}$ are the sample means of the x and y values, respectively, and $s_{p_{x,j}}$ and $s_{p_{y,j}}$ are sample standard deviations, all computed using the peaks in S_1. In the next iteration, the peaks in S_1 are removed from the replicate spectra, and peak statistics P_2 through P_m are computed, where m is a user-defined number of *prominent peaks* to include when constructing an HDC mass spectrum. The collection of peak statistics P_1 through P_m are what we define as the HDC mass spectrum. An example of the HDC mass spectrum construction process using 3 replicate mass spectra and 3 prominent peaks is shown in Fig. 2.

The primary advantage of HDC mass spectra over both traditional consensus or single measured mass spectra is the additional peak variability information captured as sample standard deviations. Using these parameters, each peak statistic in an HDC mass spectrum can be identified by a 2D probability distribution. If the identified 2D probability distribution is an accurate representation of the uncertainty associated with the peak statistics—as will be the case with increasing number of replicate mass spectra—then the similarity between any two HDC mass spectra can be approximated as a function of 2D probability distribution comparisons.

A natural place to start is by assuming peak statistics can be identified with a standard bivariate normal distribution. That is, a peak statistic $P = (\overline{p}_x, \overline{p}_y, s_{p_x}, s_{p_y})$ can be identified with the function

$$f_P(x,y) = \frac{1}{2\pi s_{p_x} s_{p_y}} e^{-\frac{1}{2}\left[\left(\frac{x-\overline{p}_x}{s_{p_x}}\right)^2 + \left(\frac{y-\overline{p}_y}{s_{p_y}}\right)^2\right]}. \quad (1)$$

To evaluate the similarity between any two peak statistics, P and Q, we compare the distributions $f_P(x, y)$ and $f_Q(x, y)$. This comparison is done by determining the cosine similarity between f_P and f_Q in the *space* $L^2(\mathbb{R}^2)$. While a discussion of function spaces is outside of the scope of this paper, we can draw an analogy to current mass spectral similarity computations to provide some context for our readers. In particular, much like how we can think of mass spectra as vectors in $\mathbb{R}^n$ (which is normally implied rather than explicitly stated), the functions f_P and f_Q are vectors in the function space $L^2(\mathbb{R}^2)$. And so we can measure the similarity of these functions as the cosine of the angle between their vectors in the *space* $L^2(\mathbb{R}^2)$.

In this space, the inner product $\langle f_P, f_Q \rangle_{L^2(\mathbb{R}^2)}$ is given by

$$\langle f_P, f_Q \rangle_{L^2(\mathbb{R}^2)} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_P(x,y) f_Q(x,y)\, dy\, dx, \quad (2)$$

and the norm $\|f_P\|_{L^2(\mathbb{R}^2)}$ in $L^2(\mathbb{R}^2)$ generated by this inner product is given

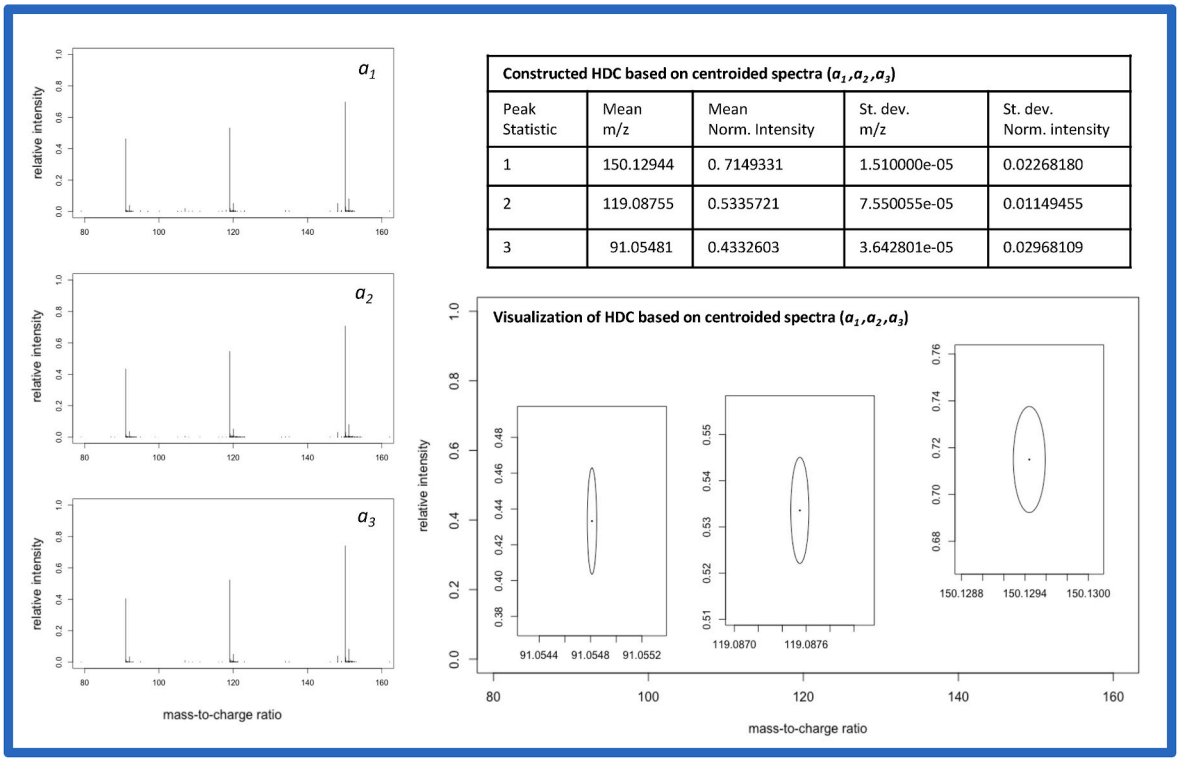

Fig. 2. An overview of how a high-dimensional consensus (HDC) mass spectrum is constructed from 3 replicate mass spectra, assuming 3 prominent peaks. Number of replicate mass spectra and prominent peaks are user-defined parameters and may be optimally selected for various applications.

by

$$\|f_P\|_{L^2(\mathbb{R}^2)} = \sqrt{\langle f_P, f_P\rangle_{L^2(\mathbb{R}^2)}} = \sqrt{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f_P(x,y)^2\ dy\ dx}, \tag{3}$$

therefore the cosine of the angle between f_P and f_Q in $L^2(\mathbb{R}^2)$ is computed by

$$\begin{aligned}\cos(\theta_{f_P,f_Q}) &= \frac{\langle f_P, f_Q\rangle_{L^2(\mathbb{R}^2)}}{\|f_P\|_{L^2(\mathbb{R}^2)}\|f_Q\|_{L^2(\mathbb{R}^2)}} \\ &= \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f_P(x,y) f_Q(x,y)\ dy\ dx}{\sqrt{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f_P(x,y)^2\ dy\ dx}\sqrt{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f_Q(x,y)^2\ dy\ dx}}.\end{aligned} \tag{4}$$

While equation (4) appears unwieldly, because of the structure of functions f_P and f_Q, surprisingly we have the exact computation

$$\cos(\theta_{f_P,f_Q}) = \sqrt{\frac{4 s_{p_x} s_{q_x} s_{p_y} s_{q_y}}{\left(s_{p_x}^2 + s_{q_x}^2\right)\left(s_{p_y}^2 + s_{q_y}^2\right)}} e^{-\frac{1}{2}\left[\left(\frac{\bar{p}_x - \bar{q}_x}{\sqrt{s_{p_x}^2 + s_{q_x}^2}}\right)^2 + \left(\frac{\bar{p}_y - \bar{q}_y}{\sqrt{s_{p_y}^2 + s_{q_y}^2}}\right)^2\right]},$$

which can be further disaggregated into its x and y components,

$$\cos(\theta_{f_P,f_Q}) = \left(\sqrt{\frac{2 s_{p_x} s_{q_x}}{s_{p_x}^2 + s_{q_x}^2}} e^{-\frac{1}{2}\left(\frac{\bar{p}_x - \bar{q}_x}{\sqrt{s_{p_x}^2 + s_{q_x}^2}}\right)^2}\right)\left(\sqrt{\frac{2 s_{p_y} s_{q_y}}{s_{p_y}^2 + s_{q_y}^2}} e^{-\frac{1}{2}\left(\frac{\bar{p}_y - \bar{q}_y}{\sqrt{s_{p_y}^2 + s_{q_y}^2}}\right)^2}\right). \tag{5}$$

Thus, for a given pair of peak statistics P and Q, we define their similarity by

$$\theta(P,Q) = \cos(\theta_{f_P,f_Q}), \tag{6}$$

where all of the terms are as defined above. The similarity measure $\theta(P, Q)$ will evaluate between 0 and 1, where 0 indicates no similarity between the peak statistics and 1 implies the peak statistics are identical.

With peak statistic similarity measure (6) in place, we can define the similarity between a pair of HDC mass spectra in the following way: Let $A = \{P_1, P_2, \ldots, P_{m_1}\}$ and $B = \{Q_1, Q_2, \ldots, Q_{m_2}\}$ be two HDC mass spectra being compared. We consider at most $m = \min(m_1, m_2)$ peak statistics in the comparison process.

To begin, we identify the peak statistic from either A or B with largest mean relative intensity. Without loss of generality, assume this is peak P_{i_1}. We then identify the peak statistic Q_{j_1} from the other HDC mass spectrum that is most similar to P_{i_1} using similarity measure θ defined in (6). This pair of peak statistics is set aside and removed from the HDC mass spectra and the process is repeated to identify pairs $\left(P_{i_2}, Q_{j_2}\right)$ through $\left(P_{i_m}, Q_{j_m}\right)$. With all of these pairs, we compute the similarity between HDC mass spectra A and B as a weighted product of peak statistics similarity using the formula

$$\varphi(A,B) = \frac{\sum_{k=1}^{m} \bar{p}_{y,i_k} \bar{q}_{y,j_k} \theta(P_{i_k}, Q_{j_k})}{\sum_{k=1}^{m} \bar{p}_{y,i_k} \bar{q}_{y,j_k}}. \tag{7}$$

The similarity score $\varphi(A, B)$ evaluates between 0 and 1, where 0 suggest that there is no similarity between HDC mass spectra and 1 means the HDC mass spectra are identical.

2.2. Experimental data

The data for our preliminary evaluation was collected with a JEOL AccuTOF LC-4G mass spectrometer (JEOL USA, Peabody, MA, USA) coupled with a DART-SVP ion source (IonSense, Saugus, MA, USA). The instrument parameters were identical to those used in the construction of the NIST DART-MS Forensic Database [18], and we collected three is-CID mass spectra (at +30 V, +60 V, and +90 V orifice 1 voltage) for each replicate using parameter switching. The DART gas temperature was set to 400 °C. Polyethylene glycol (PEG-600) was used as the mass calibration compound and a 0.1 mg/mL solution of AB-FUBINACA was used as a mass drift compensation compound. Both methamphetamine and phentermine were purchased from Cayman Chemical (Ann Arbor, MI, USA) as 1 mg/mL methanolic solutions and were diluted to an approximate concentration of 0.025 mg/mL in methanol.

3. Results and discussion

Consider the comparisons of the HDC mass spectra constructed using 15 replicate +60 V is-CID mass spectra for methamphetamine compared to itself (Fig. 3a), phentermine compared to itself (Fig. 3b), and methamphetamine compared to phentermine (Fig. 3c). Though we only display the first three peak statistic comparisons, up to 5 prominent peaks were included in the overall HDC mass spectral similarity scores presented. In the case of the first two intra-sample comparisons, the HDC mass spectra were constructed using the first 15 replicates followed by the second 15 replicates, respectively. In the third inter-sample comparison, we constructed HDC spectra using only the first 15 replicates of each compound. We can clearly see that the similarity score computed between HDC mass spectra of the same compounds (Fig. 3a and b) are higher than the score computed between HDC mass spectra constructed from different compounds (Fig. 3c). As expected, the overall HDC mass spectral similarity scores are driven by the pair-wise similarity of the dominant ions of each comparison (nominal *m/z 91*). Also, the minor peaks by which a skilled analyst might distinguish methamphetamine from phetermine (nominal *m/z 119* vs *m/z 105*) had a pair-wise similarity of 0, which is desirable. If instead of weighing pair-wise similarity by average peak intensity in Equation (7), we had selected a weighting system that highlighted the score differences in minor peaks, we could have decreased the HDC mass spectral similarity score between methamphetamine and phentermine even more. Optimally scaling parameters to improve the performance of traditional single-measurement similarity scores has been previously explored in the literature [5,19] and similar ideas may be directly applicable to HDC mass spectra.

We were surprised to see how different the observed distributions of the *m/z* 91 ion were between the HDC mass spectra of methamphetamine and phentermine. In particular, we did not expect the mean mass-to-charge ratios to vary so substantially; both compounds were measured at equal concentrations and using identical protocols. One possibility is that there was an instrument drift adding an artificial signature to the peak variability that was allowing us to so clearly distinguish the nominal *m/z* 91 ion of methamphetamine from phentermine. Our initial data collection process can be described as sequential: we collected 30 replicate mass spectra of methamphetamine with AB-FUBINACA analyzed at the beginning and after every ten measurements, and then collected 30 replicate mass spectra of phentermine with AB-FUBINACA analyzed before and after every ten measurements. And so we re-measured our compounds but alternated between which compound was measured until we had 30 replicates of each methamphetamine and phentermine. Like our first set of measurements, AB-FUBINACA was analyzed before and after every ten measurements for drift compensation. The results of comparing spectra that were collected by alternating between compounds is shown in Fig. 4.

With mass spectra measured in an alternating fashion, the mean mass-to-charge ratio of the *m/z* 91 ions between methamphetamine and phentermine were much closer together, and so the samples become more difficult to discriminate. That said, it appears that the variability in the peak intensity of the *m/z* 91 ion differs in a reproducible way between methamphetamine and phentermine. We are currently

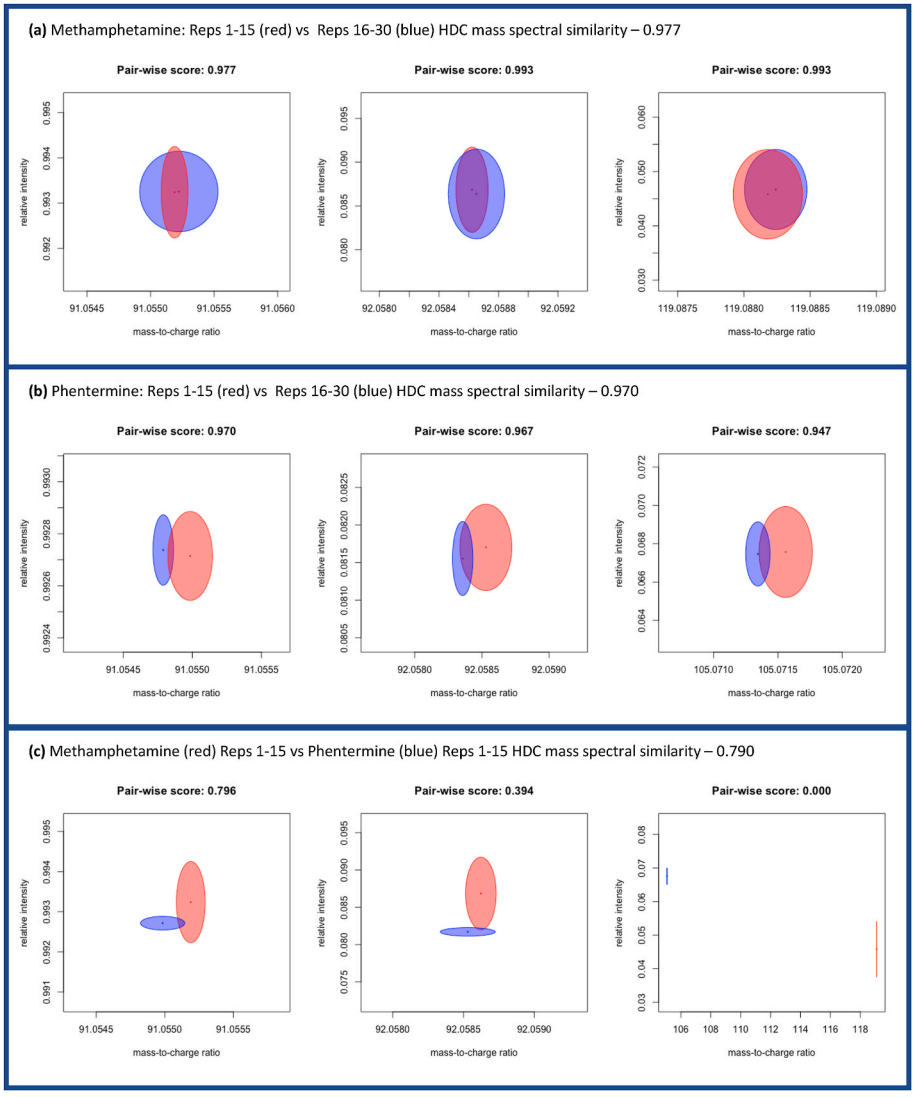

Fig. 3. Comparison of HDC mass spectra where (a) both HDC mass spectra are constructed with +60 V is-CID mass spectra of methamphetamine, (b) both HDC mass spectra are constructed with +60 V is-CID mass spectra of phentermine, and (c) one HDC mass spectrum is constructed with +60 V is-CID mass spectra of methamphetamine and the other is constructed with mass spectra of phentermine. The underlying is-CID mass spectra were collected by first measuring 30 replicates of methamphetamine followed by 30 replicates of phentermine, with AB-FUBINACA measured between every 10 replicates for drift compensation.

conducting additional experiments to determine if the discriminatory signature we observe with peak intensity variability is reproducible or whether it was introduced by the way the measurements were taken.

Because we had two sets of measurements, one where compounds were analyzed sequentially and the second where compounds were analyzed simultaneously (but alternating), we thought it might be interesting to compare how similar the spectra from each set of measurements were (Fig. 5). In this comparison, we constructed HDC mass spectra using all 30 available replicate mass spectra for each data collection method. From these results, it is clear that the observed variability of these measurements is sensitive to measurement conditions. We can see that the spectra were clearly drift-compensated in different directions, and one could easily incorrectly interpret that the sets of mass spectra are measurements of different compounds.

As noted earlier, an analyst is not likely to be analyzing pure compounds like methamphetemine and phentermine with DART-MS. And if they were looking at pure compounds, they would not be looking at single fragmentation is-CID mass spectra independently of the others that were collected simultaneously. However, we thought the example in Fig. 1 would be an interesting space to begin building a mathematical framework for incorporating measurement variability when comparing high-resolution mass spectra and we were not disappointed. We learned that there is a *signature* in the way peak variability differs between the dominant ion of the +60 V is-CID methamphetamine and phentermine (the same is true of the +90 V mass spectra but the results are not shown). We also learned that how we take measurements can affect this signature, which can subsequently affect result interpretation; with current instrumentation, it is probably best that HDC mass spectral comparisons are done with replicate mass spectra that were measured under near-identical conditions. These preliminary results have left us with several open questions:

1. What is the best way to take replicate measurements using DART-MS to minimize the likelihood of introducing an artificial signature into the observed variability? Can we control hardware and software parameters to provide more reliable measurements? There is precedent for numerical optimization of instrument parameters in mass spectrometry [20] that may be applicable to DART-MS.
2. Is there an optimal number of replicate measurements to accurately capture peak variability? While theory might suggest that we better capture peak variability with increasing number of replicates, we may also introduce deceptive variability signatures due to time-dependent sources of noise (e.g., instrument drift, ambient changes in the lab). There are also practical aspects of measurements that should be considered, and there may be diminished returns beyond a certain number of replicate measurements.
3. How will this method perform using replicate measurements from other high-resolution mass spectrometry platforms? For example,

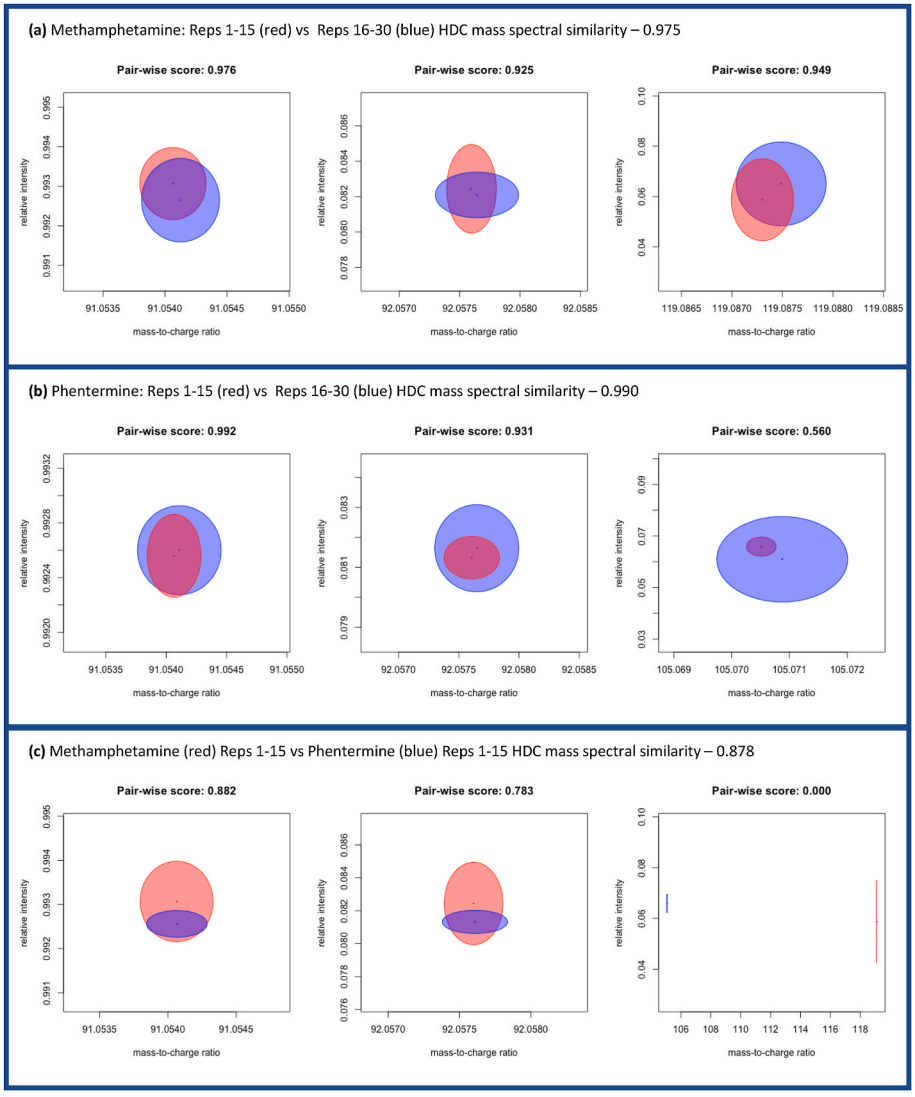

Fig. 4. Comparison of HDC mass spectra where (a) both HDC mass spectra are constructed with +60 V is-CID mass spectra of methamphetamine, (b) both HDC mass spectra are constructed with +60 V is-CID mass spectra of phentermine, and (c) one HDC mass spectrum is constructed with +60 V is-CID mass spectra of methamphetamine and the other is constructed with mass spectra of phentermine. The underlying is-CID mass spectra were collected by alternating between measurements of methamphetamine and phentermine for 30 replicates each, with AB-FUBINACA measured between every 10 replicates for drift compensation.

will chromatography-coupled high-resolution mass spectrometry provide more reliable signatures in variability? We can imagine that methods like liquid chromatography (LC)-MS greatly reduce the likelihood of contamination-based variability as compared to ambient ionization systems.

Addressing these questions, both through more highly-controlled data collection with ambient ionization platforms and by experimentation with completely different instrumentation, will allow us to identify the complete scope of application areas that might benefit from collecting multiple replicate measurements and using an HDC mass spectral similarity approach to sample discrimination.

4. Conclusion

In this paper, we introduced a new mathematical approach for computing the similarity between sets of high-resolution mass spectra. The method incorporates peak variability as observed through replicate mass spectra and thus more completely represents the similarity of two sets of data. Through example, we showed that this method can help distinguish the near-identical mass spectra of methamphetamine and phentermine when the spectra were measured in a consistent way. We also showed that if measurements were inconsistent, the introduced variability makes measurements of the same compound appear like they are different compounds. While this last result may leave some readers questioning the value of including variability when computing similarity, we are of the opinion that this added dimension of discrimination is important for making well-informed claims about the *certainty* that two sets of mass spectra are measurements of the same or different compounds.

Disclaimer

Official contribution of the National Institute of Standards and Technology (NIST); not subject to copyright in the United States. Certain commercial products are identified in order to adequately specify the procedure; this does not imply endorsement or recommendation by NIST, nor does it imply that such products are necessarily the best available for the purpose.

CRediT authorship contribution statement

Matthew J. Roberts: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition. **Arun S. Moorthy:** Conceptualization, Methodology, Software, Validation, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Funding

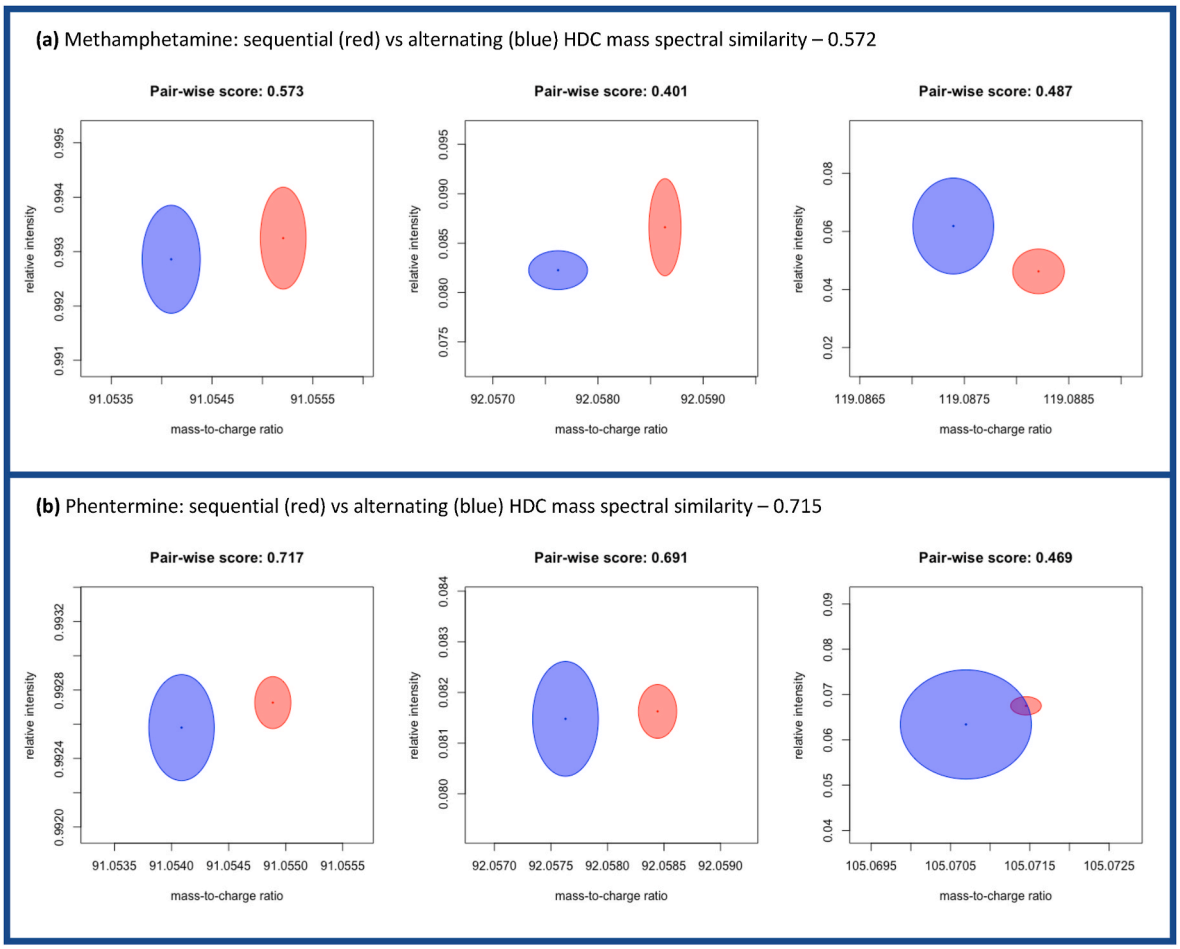

Fig. 5. Comparison of HDC mass spectra where (a) both HDC mass spectra are constructed with +60 V is-CID mass spectra of methamphetamine, and (b) both HDC mass spectra are constructed with +60 V is-CID mass spectra of phentermine. The underlying is-CID mass spectra were collected either by first measuring 30 replicates of methamphetamine followed by 30 replicates of phentermine, with AB-FUBINACA measured between every 10 replicates for drift compensation (labeled as sequential) or by alternating between measurements of methamphetamine and phentermine for 30 replicates each, with AB-FUBINACA measured between every 10 replicates for drift compensation (labeled as alternating).

acquisition. **Edward Sisco:** Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Funding acquisition. **Anthony J. Kearsley:** Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Writing – original draft, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] S.E. Stein, Mass spectral reference libraries: an ever-expanding resource for chemical identification, Anal. Chem. 84 (17) (2012) 7274–7282.
[2] H. Oberacher, M. Sasse, J.-P. Antignac, Y. Guitton, L. Debrauwer, E.L. Jamin, T. Schulze, M. Krauss, A. Covaci, N. Caballero-Casero, et al., A european proposal for quality control and quality assurance of tandem mass spectral libraries, Environ. Sci. Eur. 32 (1) (2020) 1–19.
[3] O.D. Sparkman, Evaluating electron ionization mass spectral library search results, J. Am. Soc. Mass Spectrom. 7 (4) (1996) 313–318.
[4] F. McLafferty, R. Hertel, R. Villwock, Probability based matching of mass spectra. rapid identification of specific compounds in mixtures, Org. Mass Spectrom. 9 (7) (1974) 690–702.
[5] S.E. Stein, D.R. Scott, Optimization and testing of mass spectral library search algorithms for compound identification, J. Am. Soc. Mass Spectrom. 5 (9) (1994) 859–866.
[6] K.X. Wan, I. Vidavsky, M.L. Gross, Comparing similar spectra: from similarity index to spectral contrast angle, J. Am. Soc. Mass Spectrom. 13 (1) (2002) 85–88.
[7] I. Koo, X. Zhang, S. Kim, Wavelet-and fourier-transform-based spectrum similarity approaches to compound identification in gas chromatography/mass spectrometry, Anal. Chem. 83 (14) (2011) 5631–5638.
[8] S. Nahnsen, A. Bertsch, J. Rahnenfuhrer, A. Nordheim, O. Kohlbacher, Probabilistic consensus scoring improves tandem mass spectrometry peptide identification, J. Proteome Res. 10 (8) (2011) 3332–3343.
[9] S. Kim, I. Koo, J. Jeong, S. Wu, X. Shi, X. Zhang, Compound identification using partial and semipartial correlations for gas chromatography–mass spectrometry data, Anal. Chem. 84 (15) (2012) 6477–6487.
[10] N. Garg, C.A. Kapono, Y.W. Lim, N. Koyama, M.J. Vermeij, D. Conrad, F. Rohwer, P.C. Dorrestein, Mass spectral similarity for untargeted metabolomics data analysis of complex mixtures, Int. J. Mass Spectrom. 377 (2015) 719–727.
[11] A.S. Moorthy, W.E. Wallace, A.J. Kearsley, D.V. Tchekhovskoi, S.E. Stein, Combining fragment-ion and neutral-loss matching during mass spectral library searching: a new general purpose algorithm applicable to illicit drug identification, Anal. Chem. 89 (24) (2017) 13261–13268.
[12] A.S. Moorthy, A.J. Kearsley, Pattern similarity measures applied to mass spectra, Progress in Industrial Mathematics: Success stories (2021) 43–54.
[13] A.J. Kearsley, A. Moorthy, Identifying Fentanyl with Mass Spectral Libraries, ChemRxiv, Cambridge Open Engage, Cambridge, 2021, https://doi.org/10.26434/chemrxiv.14176952.v1.
[14] Y. Li, T. Kind, J. Folz, A. Vaniya, S.S. Mehta, O. Fiehn, Spectral entropy outperforms ms/ms dot product similarity for small-molecule compound identification, Nat. Methods 18 (12) (2021) 1524–1531.

[15] X. Yang, P. Neta, S.E. Stein, Quality control for building libraries from electrospray ionization tandem mass spectra, Anal. Chem. 86 (13) (2014) 6393–6400.
[16] B.J. Place, Development of a data analysis tool to determine the measurement variability of consensus mass spectra, J. Am. Soc. Mass Spectrom. 32 (3) (2021) 707–715.
[17] X. Du, A. Smirnov, W. Jia, A preparatory study of how to construct consensus mass spectra of recurrent unknown metabolites from untargeted gc–ms metabolomics data, Int. J. Mass Spectrom. 427 (2018) 73–78.
[18] E. Sisco, A.S. Moorthy, L.M. Watt, Creation and release of an updated nist dart-ms forensics database, J. Am. Soc. Mass Spectrom. 32 (3) (2021) 685–689.
[19] S. Kim, I. Koo, X. Wei, X. Zhang, A method of finding optimal weight factors for compound identification in gas chromatography–mass spectrometry, Bioinformatics 28 (8) (2012) 1158–1163.
[20] W.E. Wallace, C.M. Guttman, K.M. Flynn, A.J. Kearsley, Numerical optimization of matrix-assisted laser desorption/ionization time-of-flight mass spectrometry: application to synthetic polymer molecular mass distribution measurement, Anal. Chim. Acta 604 (1) (2007) 62–68.