Journal of Liquid Chromatography and Gas Chromatography, Europe
https://technicalreport.org/lcgceu/

LC GC Eu., Vol. 39, No.2, 01-10, 2026

Technical Reports Press

RESEARCH ARTICLE

Collision-Energy Stability and Repository Visibility in a 1001-Compound Food-Toxicant HRMS/MS Library

Waldeck, D. H. [1*] **and Retzmann, A.**[1]

[1] *University of Pittsburgh*
[*] *Corresponding author*

Corresponding author:
dave@pitt.edu

Abstract
HRMS/MS library for screening food safety needs more than compound coverage as proof of operational worth. Robust collision-energy, unresolved workload by chemical class, visibility at repository level, and ion form–level agreement for compounds found in external libraries are other necessary factors. This research evaluates a positive-ion HRMS/MS library of food-toxicants containing 1001 compounds and 6993 manually curated spectra. The numerical information available consists of six normalized-collision-energy annotation rates, five chemical-class strata, repository-visibility numbers, and ten compound-level modified-cosine scores. NCE30 generated the highest single correct-annotation rate of 63.4%, but the use of the library would be much more justified by the use of an NCE15-NCE60 collision energy shelf having an average annotation rate of 62.73%, which is 2.63 percentage points higher than that of an NCE75-NCE90 shelf. Pesticides make up the greatest unresolved workload share of 41.2%, while organic contaminants have the greatest scarcity-adjusted vulnerability despite being just 5.3% of compounds. The comparison with public repositories shows that there are 216 compounds, 21.6% of the library, lacking in four repositories; cleaned GNPS and MS2Query planar InChIKey overlap differ in just 16 compounds. The ten-entry compound comparison has two concordant, three provisional, and five discordant entries; mean modified cosine score is 0.43, and normalized triage entropy is 0.94. These figures provide routine-use value for the NCE15-NCE60 collision energy shelf and concordant shared entries, while pointing to the need to curate pesticides, retain organic contaminant entries, and distribute 216 public-repository absent compounds.

Keywords: food safety, LC-HRMS/MS, spectral library, curation readiness, collision energy, repository visibility, modified cosine, compound triage

Received: 16 November 2025
Accepted: 20 April 2026
Available online: 30 June 2026

1. Introduction

A growing need for multiclass food safety screening calls for HRMS/MS libraries able to represent chemically varied contaminants and residues under consistent acquisition conditions. The span of chemicals included is vast: pesticide residues, veterinary drug residues, plant and fungal toxins, processing-related chemicals, industrial contaminants, and chemicals that cut across conventional classes may all find themselves included in the same screening procedure [1,2]. The reason why LC-HRMS/MS is promising in this context is that accurate precursor masses, isotope patterns, adducts, retention times, and fragment ions can be recorded in a manner that enables targeted screening, suspect screening, retrospective screening, and interlaboratory comparability [3,4]. Recently published reviews of food analysis have highlighted its significance for general surveillance of contaminants and food/feed analysis [5,6]. Food-toxicant applications have been proven by dedicated workflows for plant toxins, mycotoxins, and phytoestrogens [7]. Recent suspect and non-target studies have extended these applications to natural toxins and complex mixtures [8,9]. The latest reviews continue to stress the agricultural relevance of organic contaminants in realistic food and exposure contexts [10].

A crucial characteristic of the library is thus not merely its size. The library may include many chemicals but still lack balanced performance, due to changing collision energies, disproportionate burden carried by some classes of chemicals, lack of important compounds in the public repositories, or poor spectral agreement of the supposedly shared

Cite as: Waldeck, D. H. and Retzmann, A.. Collision-Energy Stability and Repository Visibility in a 1001-Compound Food-Toxicant HRMS/MS Library. LC GC Eu., 39(2) 2026: 01-10.

This work is licensed under Creative Commons Attribution-NonCommercial 4.0 International License

compounds [11]. Suspect and non-target studies have emphasized the necessity of communication problem created by wide coverage: the analyst has to report not only the fact of detection, but also the confidence level of that detection and associated uncertainties [12, 13]. Reliable confidence-level communication in HRMS/MS remains essential because accurate mass, formula plausibility, isotope fit, retention, adduct consistency, and fragment-ion match cannot serve as equally strong evidence [14]. The problem of over-interpretation is particularly acute in food-toxicant screening because a library result can inform subsequent review workload, confirmation, and communication of possible exposure [15].

Food-toxicant libraries operate within a larger context of the public repository system. MassBank, GNPS, MoNA, PubChem, HMDB, and chemical-classification resources allow for comparison of curated spectra and identifiers outside the generating lab [16, 17]. ENTACT mixture spectra available in MassBank and PubChem exemplify the benefits of public spectral coverage for exposomics and chemical surveillance [18]. Public visibility, however, does not amount to spectral agreement. An overlap of the planar InChIKeys can mark structural presence in another repository but does not guarantee agreement on adduct, ionization polarity, collision energy, instrumental type, purity of the precursor, and other metadata aspects [19, 20]. This is an important aspect in the case of food-toxicant chemicals since pesticide, veterinary drug, toxin, and organic contaminant entries vary in ionization efficiency and fragmentation properties. Their public representation is also influenced by broad chemical databases such as PubChem and HMDB [21]. Chemical-classification resources offer an additional framework of organization [22].

Contemporary computational mass spectrometry offers various tools for processing and comparing spectral data. MZmine and MS-DIAL are tools for reproducible LC-MS/MS feature processing and deconvolution [23, 24]. Further developments in MS-DIAL have increased its applicability for large-scale molecular profile analysis [25]. Fragmentation-tree approach and CSI:FingerID facilitate the connection between tandem spectra and molecular structures [26, 27]. SIRIUS classification expands structure-level identification of unknown metabolites [28]. MetFrag and other in silico fragmentation workflows enable candidate ranking by means of predicted fragments [29, 30]. CFM-ID allows for spectral prediction and compound identification [31, 32]. Similarity tools have also improved the reproducibility and chemical expressiveness of large-scale spectral comparison. The matchms platform unifies spectral processing and similarity evaluation [33]. Spec2Vec and MS2DeepScore introduce learning-based representations for spectral similarity [34, 35]. MS2Query and novel cross-ionization models extend the scope of analogue search in a more heterogeneous spectral space [36, 37]. Such tools enhance spectral comparison but do not eliminate the necessity of interpretation of the library evidence in order to identify the strengths and weak spots of a curated food-toxicant library. A quantitative library summary needs translation into curation priorities.

The analysis considers a positive-ion food-toxicant HRMS/MS library containing 1001 chemicals and 6993 manually curated spectra [38]. The numerical data comprises six collision-energy annotation rates, five toxicant-class strata, public-repository overlap quantities, and ten compound-level modified cosines. No spectral data is reprocessed, no compounds are assigned new identities, no acquisition step or computational identification process is introduced. The research question considered here is specific: which library values show the stability of the food-toxicant library entries and energy conditions and separate them from those requiring targeted review? This research question is different from studies presenting new sensors, new acquisition procedures, new methods of spectral topology, or new chemical balance calculations. The object of this study is the curated HRMS/MS library and the goal is to summarize its readiness pattern for library curators and food safety analysts.

The analysis proceeds along four evidence dimensions. Collision-energy values reveal whether the best annotation rate is an isolated optimum or a part of a broader robust energy region. Class-level values show whether the annotation work is difficult for the largest chemical class or for the vulnerable small classes. Repository-visibility values help distinguish the compounds absent from public repositories from the visible ones. Compound-level modified cosines classify selected shared entries as concordant, provisional, or discordant. Combining these four dimensions answers the practical question: which evidence can be used for routine curation and which should inform the allocation of targeted review, public release, metadata checking, or adduct-aware spectral reconciliation?

2. Materials and Methods

2.1. Library Values and Analytical Scope

An analysis based on the library values for a positive-ion food-toxicant HRMS/MS collection with 1001 unique compounds and 6993 curated spectra [38] was conducted. Raw spectral files were not reprocessed, peak reassignment was not attempted, and no additional chemical identities were added beyond the curated library records. All calculations are based on the stated counts, percentages, overlap values, and modified-cosine scores. The goal was to translate the stated values into a curation-readiness profile that can be viewed by library curators and food-safety analysts. The approach is appropriate for a curated library, in which the compound count, spectra count, energy-specific annotation rate, repository overlap, and similarity scores can be directly calculated.

Four numerical layers were selected. Collision-energy layer consisted of six correct-annotation percentages for NCE15, NCE30, NCE45, NCE60, NCE75, and NCE90. Class layer consisted of five strata: pesticides, veterinary drugs, natural toxins, organic contaminants, and multiple classes. Compound counts, spectrum counts, and correct-annotation percentages were used for each stratum. Repository-coverage layer contained the number of compounds not found in four public repositories and two InChIKey overlap counts (cleaned GNPS overlap and MS2Query reference overlap). Compound-conflict layer contained ten shared compounds with modified-cosine values. Four layers were retained as they answer distinct decision questions. Collision-energy values reveal the effect of acquisition settings on annotation. Class values reveal potential areas of accumulated review effort. Repository-coverage values reveal whether the library provides an expansion of the space in public repositories. Modified-cosine values reveal the spectral concordance of shared compounds.

Table 1 provides a summary of the values used in the calculations and associates each set of numbers with the corresponding curation decision. The numerical values were checked internally for consistency prior to calculations. The number of classes summed up to 1001 compounds, the number of spectra summed up to 6993 spectra. The collision energy percentages were considered globally, independent of class, since the given values were assigned to each energy rather than class. The use of the input table also explains why the analysis cannot be restricted to one success percentage. Energy values solve the question of readiness for acquisitions, class values solve the question of workload, repository values solve the question of public visibility, and modified cosine values solve the question of agreement between compounds. Since each quantity supports a different decision, the subsequent results are presented with the separate quantities rather than combined into one aggregate measure.

The tabulated values were chosen due to the library audit purposes rather than peak discovery. The raw reprocessing is highly effective when the goal is to discover new peaks, adjust preprocessing parameters, or test a different candidate-ranking algorithm. The goal here is to determine if compound counts, spectrum counts, annotation rates,

Table 1: Numerical inputs.

Data layer	Tabulated quantities	Decision role
Collision energy	Six correct-annotation rates: 62.2%, 63.4%, 62.9%, 62.4%, 61.2%, and 59.0%	Distinguishes the strongest single energy from a stable energy shelf and a weaker high-energy periphery.
Class structure	Five classes with 1001 compounds and 6993 spectra	Separates high-volume unresolved workload from small-class vulnerability.
Repository visibility	216 compounds absent from four public repositories; 680 cleaned-GNPS planar InChIKey overlaps; 696 MS2Query reference overlaps	Separates external absence from repository convergence and spectral agreement.
Compound conflict	Ten modified-cosine values ranging from 0.07 to 0.98	Assigns shared compounds to concordant, provisional, and discordant evidence states.

public repository overlap, and similarity values can serve as guides in the usage of libraries. Such choice makes the calculations easy to trace from the article tables and makes the process independent of software versions, peak-picking parameters, or proprietary acquisition methods. It also makes the analysis portable; any article about library, which includes energy-specific annotation rates, class counts, public repository overlap, and some of the similarity values, can be transformed to the same curation readiness form.

Throughout the analysis two normalizations were used. Percentages were kept as listed to make their interpretation easier while the derived measures were translated to fractions for further use in equations. Compound counts were used for class burden, since the library entry decision usually deals with the possibility to annotate the compound reliably, while the spectrum counts were used to characterize representation of a class. This way no decision weight is added to the class just because it has more replicate spectra per compound. In the current dataset the number of spectra per compound was constant within all the classes, so the qualitative information provided by the class burden (compound level) and representation (spectrum level) was the same.

2.2. Relative Collision-Energy Reliability and Shelf Detection

For each normalized collision energy e the listed correct-annotation percentage was noted as A_e. The relative reliability was calculated as

$$R_e = \frac{A_e}{A_{\max}}, \tag{1}$$

where $A_{\max}$ is the highest correct-annotation percentage among the six tested energies. The strongest single energy was defined as

$$e^* = \arg\max_e A_e. \tag{2}$$

This measure preserves the ordinary maximum-based interpretation. A second calculation was then added because a maximum alone cannot distinguish an isolated optimum from a plateau. The energy-loss value was

$$L_e = A_{\max} - A_e. \tag{3}$$

A collision energy was assigned to the annotation shelf when

$$I_e = \begin{cases} 1, & L_e \le \delta, \\ 0, & L_e > \delta, \end{cases} \tag{4}$$

where $\delta = 1.2$ percentage points. Such tolerance value is equal to the observed difference between the highest energy, NCE30, and NCE15. It was chosen beforehand without looking at the results of the shelf analysis and was made intentionally low to make it impossible to include energies that were too far from the best possible value into the shelf. The shelf retention fraction was computed as

$$S_E = \frac{\sum_e I_e}{|E|}, \tag{5}$$

with $|E| = 6$. The shelf premium was computed as

$$\Delta_E = \bar{A}_{I=1} - \bar{A}_{I=0}, \tag{6}$$

where $\bar{A}_{I=1}$ is the average percentage of correctly annotated entries among shelf energies and $\bar{A}_{I=0}$ is the average outside of the shelf.

Value of δ was not determined in order to get the shelf with four energy levels. It was related to the observed difference between the best condition and the lowest level of medium energies that were still quite close to the maximum. Such conservative choice of tolerance helps when we need a library evaluation paper since it allows avoiding an excessively wide shelf that could become just another label for almost all the tested energies. Energies outside the shelf cannot be characterized as useless. They have more decision reserve comparing to the best observed annotation condition.

2.3. Class-Level Burden and Scarcity-Adjusted Vulnerability

Since annotation decisions in library review are usually made on the level of compounds, the workloads were estimated using compound counts, not spectrum counts. The annotation deficit of class c was

$$D_c = 1 - \frac{B_c}{100}, \tag{7}$$

where B_c is the listed correct-annotation percentage for the class. The compound share was

$$P_c = \frac{N_c}{\sum_c N_c}, \tag{8}$$

where N_c is the number of compounds in class c. Workload pressure was

$$W_c = P_c D_c. \tag{9}$$

The burden share was then calculated as

$$H_c = \frac{W_c}{\sum_c W_c}. \tag{10}$$

This measure estimates how much unresolved compound-equivalent annotation work each class contributes relative to all classes.

A second axis was added to avoid allowing a large class to dominate all curation decisions. Scarcity-adjusted vulnerability was defined as

$$V_c = (1 - P_c) D_c. \tag{11}$$

The normalized value was

$$V_c' = \frac{V_c}{\max_c V_c}. \tag{12}$$

The result produces high weights for classes with significant annotation gaps and poor representation. It does not represent workload pressure and it is not meant to replace it. Rather, it responds to another question: which classes could be particularly vulnerable to evidentiary loss in the case of curation based solely on the total number of entries? A class could thus have high workload pressure, high vulnerability, high workload pressure and vulnerability, or low workload pressure and vulnerability.

The two-factor structure captures a typical trade-off in curation planning. While the largest class defines the number of entries for routine curations, a less numerous one may be more vulnerable to under-curation in the event that all decisions are made based solely on the

number of items. Thus, workload pressure and vulnerability are presented side-by-side. Workload pressure is a capacity planning index, while vulnerability is a preservation index. The distinction is critical in the food-toxicant library, since smaller classes could contain high-priority food contaminants even though they are not dominating the overall compound count.

2.4. Repository-Visibility Profile and Repository Convergence

The repository visibility was analyzed separately since a compound could have high visibility outside of the library despite its poor spectral representation or vice versa. The external absence measure was

$$X_{\text{abs}} = \frac{N_{\text{abs}}}{N_{\text{all}}}, \tag{13}$$

where $N_{\text{abs}} = 216$ and $N_{\text{all}} = 1001$. Cleaned-GNPS overlap and MS2Query reference overlap were denoted O_1 and O_2. Their convergence was

$$K_{\text{pub}} = 1 - \frac{|O_1 - O_2|}{N_{\text{all}}}. \tag{14}$$

A high value suggests that the two overlap counts are similar numbers at the library level. This does not guarantee that each of the two shared spectra is valid since spectral similarity depends on the fragment-ion pattern, adduct state, and acquisition conditions.

The use of Planar InChIKey overlap for validation was conservative. It was an effective indicator of structural visibility since it grouped the records which have the same connectivity; however, it did not account for stereochemistry, acquisition method, collision energy, adduct state, and spectral quality. That is why the overlap counts were considered as an indication of repository visibility and not as a validation of spectral similarity. The following modified-cosine similarity was required due to the discrepancy between structural visibility and spectral agreement.

2.5. Compound-Level Triage and Conflict Entropy

Modified-cosine value M_i for the compound i in the ten-compound conflict subset was estimated. Disagreement was defined as

$$C_i = 1 - M_i. \tag{15}$$

Compounds with $M_i \geq 0.70$ were considered as concordant. Compounds with $0.50 \leq M_i < 0.70$ were considered as provisional. Compounds with $M_i < 0.50$ were considered as discordant. These cutoff values were considered as a triage criteria and not as an identification criteria. The goal of their use was to separate compounds by the level of spectral agreement (high, intermediate, and low).

Category dispersion was characterized by the normalized entropy

$$E_T = -\frac{\sum_g p_g \ln(p_g)}{\ln(G)}, \tag{16}$$

where p_g is the probability of compounds in category g, and $G = 3$ is the number of triage categories. Zero value suggests concentration in one evidence category, while one value implies dispersion in all three categories. Entropy was estimated in addition to the mean modified-cosine value to characterize the reason behind low mean value as uniformly weak spectra or mixture of concordant, provisional, and discordant spectra.

3. Results and Discussion

3.1. The Library Values Form a Coherent Curation Dataset

The library values formed a coherent four-layer dataset suitable for curation-readiness analysis. The 1001 compounds were allocated into 454 pesticides, 253 veterinary drugs, 206 natural toxins, 53 organic contaminants, and 35 multiple-class entries. The corresponding spectrum numbers were 3174, 1764, 1439, 371, and 245. The spectrum density was surprisingly stable across all five classes from 6.97 to 7.00 spectra per compound. The spectrum density stability proves that differences in class annotation reserves did not stem from the lack of spectra per compound. On the contrary, the decision reserve was formed by class annotation percentages, collision energy behavior, repository visibility, and compound level spectral correspondence.

The class-level correct annotation percentages were 66.7%, 60.7%, 59.7%, 56.6%, and 71.4% for pesticides, veterinary drugs, natural toxins, organic contaminants, and multiple-class entries respectively. It is immediately apparent from the data above that there is no direct relationship between the number of compounds per class and the corresponding annotation rate. The multiple-class category had the largest annotation percentage but the smallest compound count. Pesticides had the largest count and the second highest annotation percentage. Organic contaminants had the smallest annotation percentage but only a minor representation in the library. Based only on the annotation percentage, the library should be focused on the organic contaminants, while only pesticides would get priority based on the count.

The repository visibility layer also allowed for separate interpretation. The library included 216 compounds missing in four public repositories, while the clean GNPS planar InChIKey overlap covered 680 compounds and the MS2Query reference overlap included 696 compounds. The absence count clearly states that the library offers chemical coverage exceeding that of the major public repositories. Similarity in the two overlap counts suggests that visibility is relatively convergent among repositories. The ten-compound conflict subset shows that visibility does not necessarily correlate with spectral correspondence. This separation is consistent with the broader practice of open spectral repositories utilization, where library searches, molecular networking, and analogues search provide for better prioritization but are still insufficient for evidence evaluation [19, 36].

The result presented in Figure 1 provides the visual rationale for class-level interpretation. Pesticides prevail both in terms of the number of compounds and the number of spectra, yet all five strata feature approximately seven spectra per compound. Such uniformity of spectral density implies that the class-level difference in annotation workload and vulnerability should be attributed to the behavior of annotation process and class size, but not the difference in spectral availability.

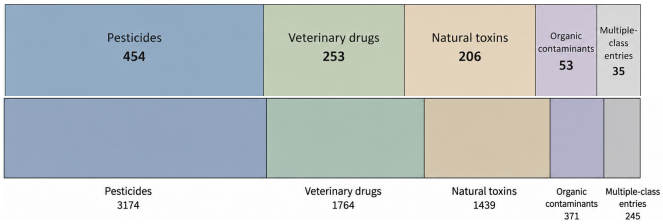

Figure 1: Class and spectra ledger.

Spectral density stability is important as it excludes any straightforward explanation of the class-level differences. If any class was characterized by a much smaller number of spectra per compound, poor annotation performance could be interpreted primarily through weaker replication. In reality, however, all classes were characterized by the same approximate value of spectral density. Differences in curation-readiness, hence, have some substance behind them and include such factors as chemical diversity inside a class, fragmentation behavior, representation in major public repositories, and the extent to which the information provided by spectra collected under different collision energies supports the same annotation decision. The observation makes the class-level interpretation even stronger as it implies that the class-level results are not an artifact of the inequality in spectrum

availability.

Value set in the table includes an important balance between breadth and depth. Breadth is represented by 1001 compounds in several food-toxicant categories and by 216 compounds lacking representation in major public repositories. Depth is represented by multiple spectra per compound and collision-energy-dependent annotation rates. Library can fail as a practical decision tool both if it features bread without depth (many entries will lack substantial support) and if it features depth without breadth (well-characterized compounds will cover only a limited part of monitoring space). Both components are available in the library value set and their balance is evaluated in the curation-readiness analysis.

3.2. Collision-Energy Behavior is Better Described as a Shelf than as a Single Optimum

Condition NCE30 proved to be the best in terms of annotation performance among all conditions used in the study. The correct-annotation rate of 63.4% and the relative reliability coefficient of 1.000 confirm this conclusion. The next-best condition, NCE45, provided 62.9% correct annotations and R_e equal to 0.992. NCE60 and NCE15 also proved to be relatively strong with 62.4%/R_e =0.984 and 62.2%/R_e =0.981 correspondingly. Conditions NCE75 and NCE90, however, showed poor performance of 61.2%/R_e =0.965 and 59.0%/R_e =0.931 correspondingly. The maximum-based interpretation of the results is thus straightforward: condition NCE30 is the most effective in the library table. Shelf-based interpretation provides more useful information as it illustrates that the most effective condition belongs to a broad and stable area.

The numbers in Table 2 indicate that NCE15, NCE30, NCE45, and NCE60 passed the shelf criterion of 1.2 percentage points. Hence, the shelf-retention ratio is 4/6, or 0.667. Inside the shelf, the mean annotation rate was 62.73%, while it was 60.10% outside the shelf, making the shelf premium 2.63 percentage points. This finding changes the conclusion about the optimal configuration for annotation decisions. If there is a narrow maximum, one should give preference to one acquisition condition. Shelf means that nearby collision energies comprise a stable decision zone and can be considered as evidence corroborating each other. As fragmentation intensity and product-ion distribution depend on collision energy, the interpretation of the shelf result is especially valuable in case of the compound with borderline spectral agreement at one energy and strong support at another one.

Table 2: Energy reliability.

Collision energy	Correct annotation (%)	R_e	L_e	Shelf status
NCE15	62.2	0.981	1.2	In shelf
NCE30	63.4	1.000	0.0	In shelf
NCE45	62.9	0.992	0.5	In shelf
NCE60	62.4	0.984	1.0	In shelf
NCE75	61.2	0.965	2.2	Outside shelf
NCE90	59.0	0.931	4.4	Outside shelf

The finding demonstrated in Figure 2 takes into account not absolute annotation rate value, but its distance from the highest rate. Such representation makes the visual impression of the high-energy peripheral region very clear. NCE90 has a gap of 4.4 percentage points from the shelf, which is much larger than the distances between all four shelf energies. The practical implication of this finding is not to discard high-energy spectra. They can be useful for certain structures. What the finding means is that the routine decision support is strongest at the NCE15–NCE60 shelf, while NCE75 and NCE90 should be considered as additional or diagnostic energies, not routine annotation conditions. This recommendation follows the general idea that tandem spectra should be analyzed considering acquisition conditions and not just as a fingerprint [39, 40].

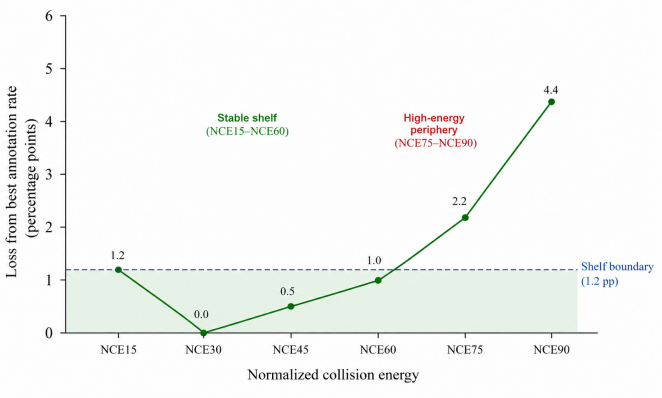

Figure 2: Collision-energy loss profile.

The close values of the four shelf energies have also a particular significance for how spectra should be considered in a screening procedure. If a compound has the same level of evidence at several shelf energies, the analyst may conclude that the annotation is robust because it is based on the consistent evidence. If the evidence appears only at the high-energy peripheral region, the finding may be still valuable, but should be treated as energy-dependent. This is particularly relevant for retrospective screening, when a laboratory compares its new spectra with library spectra acquired in different conditions. The shelf-type library entry has higher tolerance for the mismatch than an entry with the narrowly defined optimum of energies.

The 2.63-percentage-point shelf premium seems small, but it is meaningful because it describes the trend at a range of energies rather than pairwise difference. The drop from NCE30 to NCE90 is 4.4 percentage points, while differences between NCE15, NCE30, NCE45, and NCE60 are negligible. In a library with 1001 compounds, even several percentage points are equivalent to a significant number of annotation results. Thus, the finding implies a rule: medium-energy spectra should be kept as the routine comparison core, while high-energy spectra should be used for clarifying selected cases, not defining the general behavior of the library.

3.3. Class Workload and Vulnerability Separate two Curation Priorities

Class-level workload estimate gives 366.64 unresolved compound-equivalent annotations in the five-class table. For pesticides, the value is 151.18, for veterinary drugs—99.43, for natural toxins—83.02, for organic contaminants—23.00, and for multiple-class entries—10.01. Normalizing to the total unresolved workload, we get 41.2% for pesticides, 27.1% for veterinary drugs, 22.6% for natural toxins, 6.3% for organic contaminants, and 2.7% for multiple-class entries. These numbers show that pesticides are the main workload class although their annotation percentage is not the smallest. The reason for it is the size.

The values in Table 3 highlight the need for a second class dimension. The vulnerability score inversely ranks the small classes in terms of practical significance. Organic contaminants had the largest normalized vulnerability, 1.000, as this class was comprised of the greatest deficit of entries relative to its representation. Natural toxins came in second place, 0.779, and veterinary drugs in third, 0.714. Pesticides had the largest workload but the smallest normalized vulnerability among the single classes. This is not contradictory. Rather, it shows two different curation tasks. In the case that the laboratory wants to allocate review time for the most unresolved entries, pesticides are the class to target first. If the laboratory is concerned about underrepresented classes whose evidence could get lost due to volume-based

review processes, organic contaminants should be a priority.

Table 3: Class burden.

Class	Compounds	Spectra	Annotation (%)	Unresolved eq.	Burden share	V_c'
Pesticides	454	3174	66.7	151.18	0.412	0.443
Veterinary drugs	253	1764	60.7	99.43	0.271	0.714
Natural toxins	206	1439	59.7	83.02	0.226	0.779
Organic contaminants	53	371	56.6	23.00	0.063	1.000
Multiple-class entries	35	245	71.4	10.01	0.027	0.671

The two-dimensional plane in Figure 3 offers a concise interpretation of these two priorities. Pesticides are far to the right because they represent the class with the largest unresolved workload. Organic contaminants are up because of their normalized vulnerability being the largest. Veterinary drugs and natural toxins are within the middle decision space: both have a high workload and a noticeable vulnerability. Multiple-class has low workload because there are just 35 compounds within it, however, the vulnerability value is still noticeable since even a small class becomes difficult to interpret when entries span the boundaries between classes. Such an interpretation can be valuable for library management as it lets review planning take into account both throughput and preservation considerations. Without this separation, the curator could spend too much time on the largest class and pay insufficient attention to the small classes, which have low count but low quality of evidence.

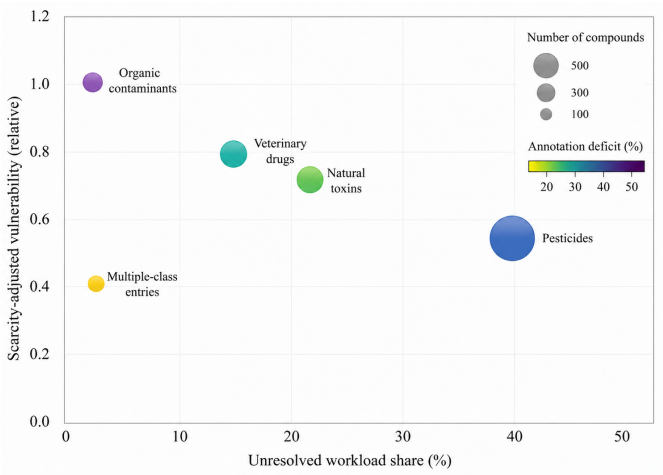

Figure 3: Class workload-vulnerability plane.

The class result helps understand the practice of food safety screening. Pesticides are usually dominating multiclass residue lists because they are generated through agriculture and are regulated. The veterinary drugs and natural toxins add different fragmentation and matrix effects. Organic contaminants are chemically diverse and could have lower reference coverage. A global annotation percent cannot show these distinctions. A curation-readiness map will not help solving all the class-specific analytical issues but will give an explicit starting point for the curator to choose how to add extra spectra, orthogonal confirmation, retention index information or manual reviews.

3.4. Repository Visibility and Spectral Agreement Remain Separate Evidence States

From the repository-visibility analysis, we concluded that 216 of 1001 compounds were not found in any of the four public repositories. This yielded an external-absence measurement of 0.216. Cleaned-GNPS planar InChIKey overlap covered 680 compounds, or 67.9% of the library. MS2Query reference overlap covered 696 compounds, or 69.5%. The two overlap percentages differ by 16 compounds, which corresponds to 1.6% of the library and the repository convergence measurement is 0.984.

The values provided in Table 4 allow for the following observations. Firstly, the library adds significant visibility since one fifth of its compounds could not be found in the considered public repositories. However, being absent in public repositories does not imply that the compounds are analytically more important than those which have already gained visibility. Secondly, cleaned-GNPS and MS2Query reference overlap counts are numerically very similar. Therefore, the fact that the spectral disagreement exists among shared compounds cannot be considered merely as an indicator of an incoherent discrepancy between public overlap counts. It should be assessed at the levels of spectrum, adducts, and acquisition conditions.

Table 4: Repository visibility.

Quantity	Count	Fraction of library
Absent from four public repositories	216	0.216
Planar InChIKey overlap with cleaned GNPS	680	0.679
Planar InChIKey overlap with MS2Query reference collection	696	0.695
Difference between overlap counts	16	0.016
Repository-convergence value K_{pub}	–	0.984

The repository ruler shown in Figure 4 puts public absence and public overlaps on the same scale. Absence can be seen as being numerically substantial, yet still smaller compared to both measures of public overlap. As such, the practical conclusion is that the library curation process needs to address two distinct issues. The first is expanding coverage for those compounds that are not present in major public repositories. The second is spectral matching for those compounds that are present in major public repositories but fail to provide a strong spectral match. Such issues need different sets of evidence. While coverage expansion is mainly about correct metadata and accessible spectra, spectral matching is more about adduct-aware matching and energy-aware comparisons with peak-by-peak analysis and possibly alternative acquisition conditions. Combining both issues into one will only make it hard to understand why a particular compound cannot be matched.

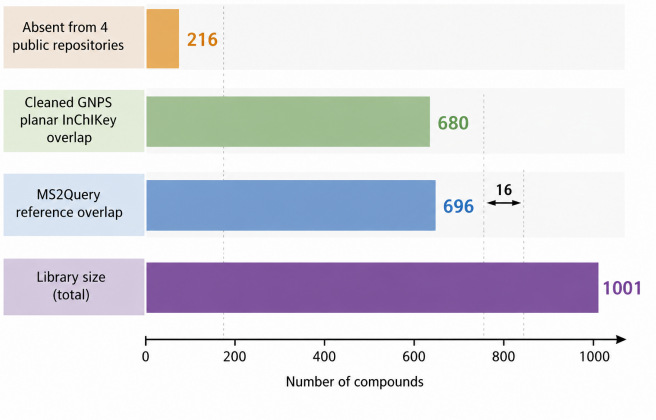

Figure 4: Repository visibility ruler.

3.5. Compound-Level Conflict Demonstrates wide Dispersion Across Triage States

The set of ten-compounds selected according to modified cosine included two concordant compounds, three provisional compounds, and five discordant compounds. The mean modified cosine was 0.43, the median was 0.415, the standard deviation was 0.31, and the range was 0.91. It is easy to see that these descriptive statistics demonstrate significant heterogeneity. However, the mean can be considered as

indicative of weak agreement only in combination with the category counts.

The compound-specific lane map in Figure 5 confines the conflict analysis to the compound level. Phorate and sulfadoxine reside in the concordant region, while bixafen, coniine, and phorate sulfoxide are located in the provisional region. The five remaining entries are located in the discordant region. The adduct-specific entries are still present since the repository overlap cannot prove that the same ion form was considered in two spectra.

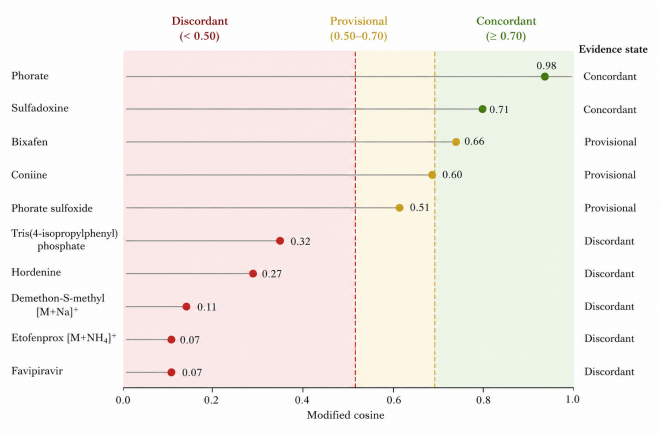

Figure 5: Compound agreement lanes.

The compound-specific triage table (Table 5) shows that presence of the common compound in two repositories does not ensure the spectral comparability. Phorate and sulfadoxine showed the modified cosine values above the threshold. Bixafen, coniine, and phorate sulfoxide are provisional and should be interpreted with supporting evidence instead of being rejected. The remaining five compounds belong to the discordant category with three values at or below 0.11. There are two discordant pesticide entries that involve the non-protonated adducts. Thus, the table proves the necessity of the adduct-specific triage approach. Presence of the common compound in the two repositories cannot serve as an indicator of its secureness.

Table 5: Compound triage.

Compound	Class assignment	Modified cosine	Decision state
Phorate	Pesticide	0.98	Concordant
Sulfadoxine	Veterinary drug	0.71	Concordant
Bixafen	Pesticide	0.66	Provisional
Coniine	Natural toxin	0.60	Provisional
Phorate sulfoxide	Pesticide	0.51	Provisional
Tris(4-isopropylphenyl) phosphate	Organic contaminant	0.32	Discordant
Hordenine	Natural toxin	0.27	Discordant
Demethon-S-methyl $[M + Na]^+$	Pesticide	0.11	Discordant
Etofenprox $[M + NH_4]^+$	Pesticide	0.07	Discordant
Favipiravir	Veterinary drug	0.07	Discordant

The ternary plot in Figure 6 clarifies the high value of the normalized entropy. The conflict subset was not concentrated in one category. In case all ten compounds were discordant, the mean would be low and entropy would be low too. In case all compounds were concordant, the curation and interpretation would be simple. It is more informative to observe several evidence states in the subset simultaneously. The plot helps to establish the tiered review process. Concordant entries can be used in routine applications as long as other evidence supports them. Provisional entries should be checked according to the retention time, precursor form and energy-specific fragment ions. Discordant entries require a more prudent treatment and should be prioritized for spectrum-level reconciliation. Entropy is helpful since it prevents the curation problem from being defined by the low mean score only.

The compound list also confirms the necessity to carry out the conflict

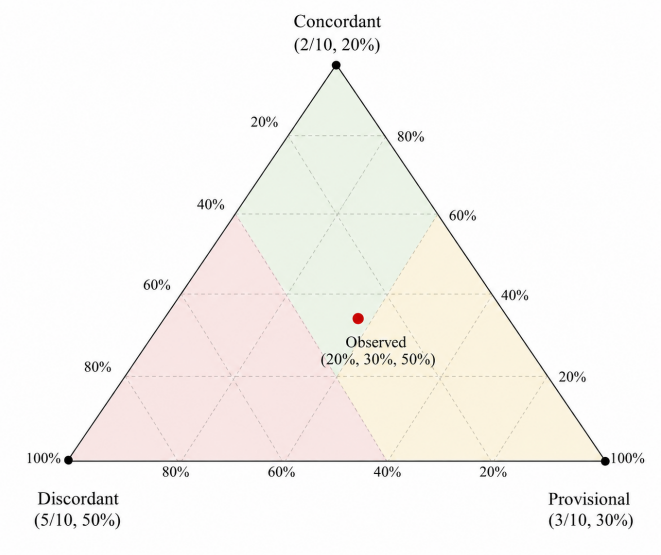

Figure 6: Triage dispersion triangle.

review on the compound-specific level. Phorate and phorate sulfoxide are chemically related entries but belong to different triage categories in the subset. Phorate is concordant, while phorate sulfoxide is provisional. Such difference warns about the inappropriateness of transfer of confidence between two related compounds without inspection of the actual spectrum. Similarly, the discordant examples contain various chemical classes and adduct forms indicating that the weak correspondence is not limited to the particular chemical class. Therefore, the triage table can be used as a diagnostic checklist: check adduct, match energy, compare major fragments, and decide whether the public spectrum and the library spectrum represent the same analytical form. This compound-specific interpretation of the table is also consistent with recent similarity literature: learned and large-scale spectral comparison can facilitate triage, but the interpretation is still dependent on ion form and metadata context.

Entropy becomes valuable since it describes the structure of the conflict problem [41]. A mean modified cosine of 0.43 can be formed by the uniformly intermediate matches, a few excellent and many poor matches, or the heterogeneous structure like in our example. The three scenarios would require different actions. The entropy value of 0.94 indicates the heterogeneity of evidence states, therefore, a unified approach would be inefficient. Concordant compounds do not need the same actions as discordant compounds. Provisional compounds cannot be treated as failures in case there is some supporting evidence to interpret them.

3.6. The Integrated Evidence Profile Enhances Library Interpretation

The integrated curation readiness map alters the library interpretation. Collision energy analysis indicates that NCE30 is the best single collision energy, whereas the NCE15-NCE60 shelf constitutes the practical evidence zone. Class analysis shows that pesticides make up the dominant class and the vulnerable organic contaminants have to be protected. Repository-visibility analysis reveals the presence of both external absence and repository convergence. Compound conflict analysis identifies shared compounds in multiple evidence states. These four insights are independent but complementary. None can be deduced from library size alone.

The new approach also reveals the priorities of further curation. In terms of collision energy, the priority is not to increase the number of spectra at all possible energies, but to stabilize the energies on

the shelf, which ensure stable routine annotation and treat the high-energy periphery as a diagnostic rather than dominating channel. In terms of class curation, the priority is to devote the standard review time to pesticides, veterinary drugs, and natural toxins, since they make up most of the unresolved compound-equivalent workload and simultaneously protect organic contaminants from being under-curated. In terms of repository visibility, the priority is to expand metadata of the 216 absent compounds, while simultaneously solving weak match problems of the already visible ones. Finally, in terms of compound conflict, the priority is to perform triaging.

Figure 7 shows the curation action matrix that converts four evidence layers into library management decisions. Energy-shelf evidence suggests routine use of the NCE15-NCE60 region, class workload directs review efforts to pesticides, small-class vulnerability protects organic contaminants, public absence directs metadata expansion of the 216 absent compounds, repository convergence makes cross-resource checking necessary, and compound conflict directs manual reconciliation of provisional and discordant entries.

Evidence layer	Routine use	Review allocation	Metadata expansion	Spectrum reconciliation
Collision-energy shelf	Apply NCE15–60 as preferred energy range	Inspect NCE75–90 records systematically	Record full energy series and performance indicators	Resolve energy-specific annotation disagreements
Class workload	Prioritize routine use for lower workload classes	Allocate most review effort to pesticides	Tag unresolved burden per class	Reconcile class-level annotation inconsistencies
Small-class vulnerability	Maintain coverage of vulnerable small classes	Safeguard organic contaminant entries	Add saarcity and vulnerability flags	Improve spectra quality for small classes
Repository absence	Use internal library with confidence	Document and monitor absent compounds	Add repository absence annotations	Generate spectra for absent compounds
Repository convergence	Leverage convergent overlap for confidence	Monitor divergence between resources	Record source identifiers and linkages	Address conflicting annotations across resources
Spectral conflict	Accept concordant matches	Manually inspect provisional and discordant cases	Add similarity scores and adduct information	Curate spectra for discordant entries

Figure 7: Curation action matrix.

This decision-oriented interpretation is coherent with the broader development trends in computational mass spectrometry. Modern databases and community-based molecular networking systems broaden the range of spectra that can be compared and shared [19, 42]. Standardized similarity processing and learned spectra representations extend these capabilities even further [33, 34]. Deep-learning and analogue-based systems add scalable pattern-aware matching [35, 36]. Prediction of cross-ionization states and rapid search extensions continue to enhance this analytical toolkit [37, 43]. None of these developments replaces the necessity of accounting for the evidence quality. Large-scale and pattern-aware search increases the reach, but also the need for clear filters, confidence tiers, and curation priorities [44, 45]. Curation readiness profile creates those priorities using the library values that are immediately visible: counts, percentages, overlap amounts, and modified cosine scores.

The analysis has boundaries. It does not replace the spectral reanalysis, chromatographic validation, matrix-specific confirmation, ion ratio evaluation, and regulation-based confirmation workflow. Also, it does not suggest that a modified cosine threshold is universal for all instruments, adducts, and compound classes. The thresholds used here were triage thresholds, and not all of them may be universal. Still, the present analysis is valuable in that it extracts the additional decision values from the tabulated 1001-compound library values.

Four evidence layers may be seen as a series of increasingly specific questions. Energy layer asks for which acquisition conditions the most stable evidence is available. Class layer asks about the workload accumulation and small classes that need protection. Public layer asks whether the library adds the external chemical visibility or simply duplicates it. Compound layer asks about the agreement of the shared entries. The progression through those layers prevents premature conclusions: the class may look acceptable in terms of percentage, but actually make up the dominant workload. The compound may be visible in the public repository, but still be discordant. The collision energy may be the maximum, but belong to the energy shelf. These are exactly the differences that the curation readiness map preserves.

From the perspective of practical library governance, the analysis suggests the staged review strategy. The first stage is energy stabilization: preserve and document the NCE15-NCE60 shelf because it contains the most stable routine evidence. The second stage is class workload review: assign the review time to pesticides, veterinary drugs, and natural toxins because they dominate the unresolved compound equivalent workload. The third stage is the small class protection: audit the organic contaminants so that the weakest and smallest class will not get lost in volume-based workflow. The fourth stage is the repository-visibility and conflict review: separate the compounds absent from public repositories from compounds that are publicly visible but spectrally discordant. This staged policy is simple enough to be implemented for future library releases, yet grounded in the numerical values of the library record.

The organization of four evidence layers also enhances the interpretability, since each layer corresponds to a different action, a principle that is consistent with the spectral network and annotation studies that preserve the chemical relationships rather than flatten them into single score [46, 47]. Energy-shelf evidence guides the acquisition and matching policy. Class-burden evidence guides the distribution of routine review effort. Repository-visibility evidence guides the dissemination and interoperability priorities. Conflict-entropy evidence guides the compound-level reconciliation. This action-oriented structure is preferable to the aggregate score because the same aggregate value would hide whether the library weakness arises from energy instability, class imbalance, public absence, or spectral disagreement. The same value could therefore imply the very different curatorial response. By separating the layers, the analysis allows the curator to choose the proper way to handle the problem: acquiring or preserving spectra at particular collision energies, expanding metadata of the underrepresented public compounds, auditing the vulnerable class, or re-checking the discordant adduct state. This approach also facilitates the transparent communication with the library users: a laboratory can tell whether the compound belongs to the well-supported energy shelf, but still requires caution because it is absent from the public repositories, or whether the compound is visible in the public resources, but spectrally discordant. Such statements are more informative than the single declaration of match or non-match, and are consistent with the graded confidence communication in high resolution mass spectrometry.

Figure 8: Integrated evidence board.

The final evidence board in Figure 8 gives the final synthesis of the analysis results. Four cards reveal that the library has a stable routine energy shelf, concentrated pesticide review burden, significant public coverage reserve, and a dispersed conflict structure. The figure is compact on purpose, since the final decision is not to collapse the

analysis into one score, but to preserve the four curation signals.

4. Conclusion

The research question examined whether the values of a 1001-compound food-toxicant HRMS/MS library, listed above, could determine which part of the library can be used routinely for curation and which requires targeted review without adding a new acquisition and identification step to the process. The answer is yes, as long as the notion of readiness is understood as an independent evidence dimension rather than an overall percentage value. This library shows a clear readiness in NCE15–NCE60 collision-energy shelf where the average annotation rate was 62.73%, and the loss from the maximum NCE30 value was 1.2 percentage points. The conclusion means that the practical applicability of the energy spectrum of this library is not defined by a single optimum but spread in a range of four neighboring energy settings.

The same pattern allows identifying the areas where routine use of the library requires extra caution. The pesticides should be prioritized for the review because they constitute 41.2% of unresolved compound-equivalent entries. The organic contaminants should be considered separately since they have a smaller absolute number (53 entries), but their adjusted sensitivity level is the highest in this group of five toxicant classes. Repository analysis introduces a third decision criterion: 216 compounds missing in four public repositories represent the public-visibility advantage of this library, while the overlaps with GNPS (680) and MS2Query (696) databases show sufficient visibility of library entries at the identifier level. The difference between the two overlap figures is 16 compounds – small enough in absolute terms of this library but not enough to guarantee the spectral equivalence. Compound-level comparison produces a final answer: 2 concordant, 3 provisional, and 5 discordant shared compounds, thus resulting in the average modified cosine of 0.43 and the normalized triage entropy of 0.94. The subset of conflicts is heterogeneous and cannot be resolved by a single accept/reject rule.

The practical conclusion is that the current library is ready to be communicated as the profile of readiness. Routine annotation is feasible in NCE15-NCE60 shelf and concordant compound matching; the effort should be focused on pesticide review, protection of smaller organic-contaminant class, public-visibility contribution of the 216 compounds absent from public repositories, and adduct-aware resolution of the 5 discordant shared spectra. This conclusion applies specifically to this library containing 1001 compounds, 6993 spectra, six collision energies, five toxicant classes, three repository visibility measures, and ten modified cosine conflicts. The analysis thus does not make the assumption that there is a single standard for the quality of the library. Instead, what can be said is that the usefulness of an HRMS/MS library is greatest when the strengths and weaknesses of the library are reported as evidence states separately.

References

[1] Kaufmann, A. (2012). The current role of high-resolution mass spectrometry in food analysis. Analytical and bioanalytical chemistry, 403(5), 1233-1249.

[2] Lehotay, S. J., & Chen, Y. (2018). Hits and misses in research trends to monitor contaminants in foods. Analytical and Bioanalytical Chemistry, 410(22), 5331-5351.

[3] Lehotay, S. J. (2024). Food safety analysis 2.0. Analytical and Bioanalytical Chemistry, 416(3), 609-610.

[4] Hernández-Mesa, M., & Moreno-González, D. (2022). Current role of mass spectrometry in the determination of pesticide residues in food. Separations, 9(6), 148.

[5] Sun, Q., Dong, Y., Wen, X., Zhang, X., Hou, S., Zhao, W., & Yin, D. (2023). A review on recent advances in mass spectrometry analysis of harmful contaminants in food. Frontiers in Nutrition, 10, 1244459.

[6] Gavage, M., Delahaut, P., & Gillard, N. (2021). Suitability of high-resolution mass spectrometry for routine analysis of small molecules in food, feed and water for safety and authenticity purposes: A review. Foods, 10(3), 601.

[7] Bessaire, T., Ernest, M., Christinat, N., Carrères, B., Panchaud, A., & Badoud, F. (2021). High resolution mass spectrometry workflow for the analysis of food contaminants: Application to plant toxins, mycotoxins and phytoestrogens in plant-based ingredients. Food Additives & Contaminants: Part A, 38(6), 978-996.

[8] Zhang, Y., Chen, T., Wang, Z., Liang, W., Wang, X., Zhang, X., ... & Xu, G. (2025). High-resolution mass spectrometry-based suspect and nontarget screening of natural toxins in foodstuffs and risk assessment of dietary exposure. Environmental Pollution, 365, 125338.

[9] Motteau, S., Dervilly, G., Cariou, R., Margalef, M., Lamoree, M., Hamers, T., ... & Antignac, J. P. (2025). Determination of Chemical Mixtures in Environmental, Food, and Human Samples Using High-Resolution Mass Spectrometry-Based Suspect Screening Approaches. Environmental Science & Technology, 59(39), 21265-21277.

[10] Vergara-Luis, I., Righetti, L., Olivares, M., Padilla-Gonzalez, F., Prieto, A., Arrizabalaga-Larrañaga, A., & Ayala-Cabrera, J. F. (2025). A review of suspect and non-target screening of agriculture-related organic contaminants. Analytica Chimica Acta, 344774.

[11] González-Gaya, B., Lopez-Herguedas, N., Bilbao, D., Mijangos, L., Iker, A. M., Etxebarria, N., ... & Zuloaga, O. (2021). Suspect and non-target screening: the last frontier in environmental analysis. Analytical Methods, 13(16), 1876-1904.

[12] Pourchet, M., Debrauwer, L., Klanova, J., Price, E. J., Covaci, A., Caballero-Casero, N., ... & Antignac, J. P. (2020). Suspect and non-targeted screening of chemicals of emerging concern for human biomonitoring, environmental health studies and support to risk assessment: From promises to challenges and harmonisation issues. Environment international, 139, 105545.

[13] Hollender, J., Schymanski, E. L., Ahrens, L., Alygizakis, N., Béen, F., Bijlsma, L., ... & Krauss, M. (2023). NORMAN guidance on suspect and non-target screening in environmental monitoring. Environmental Sciences Europe, 35(1), 75.

[14] Summer, L. W., Amberg, A., Barrett, D., Beale, M. H., Beger, R., Daykin, C. A., ... & Hankemeier, T. (2007). Proposed minimum reporting standards for chemical analysis. Metabolomics, 3(3), 211-221.

[15] Kind, T., & Fiehn, O. (2007). Seven Golden Rules for heuristic filtering of molecular formulas obtained by accurate mass spectrometry. BMC bioinformatics, 8(1), 105.

[16] Horai, H., Arita, M., Kanaya, S., Nihei, Y., Ikeda, T., Suwa, K., ... & Nishioka, T. (2010). MassBank: a public repository for sharing mass spectral data for life sciences. Journal of mass spectrometry, 45(7), 703-714.

[17] Neumann, S., Meier, R., Wenk, M., Elapavalore, A., Nishioka, T., Schulze, T., ... & Schymanski, E. L. (2026). MassBank: an open and FAIR mass spectral data resource. Nucleic Acids Research, 54(D1), D601-D606.

[18] Elapavalore, A., Kondić, T., Singh, R. R., Shoemaker, B. A., Thiessen, P. A., Zhang, J., ... & Schymanski, E. L. (2023). Adding open spectral data to MassBank and PubChem using open source tools to support non-targeted exposomics of mixtures. Environmental Science: Processes & Impacts, 25(11), 1788-1801.

[19] Nothias, L. F., Petras, D., Schmid, R., Dührkop, K., Rainer, J., Sarvepalli, A., ... & Dorrestein, P. C. (2020). Feature-based molecular networking in the GNPS analysis environment. Nature methods, 17(9), 905-908.

[20] Aron, A. T., Gentry, E. C., McPhail, K. L., Nothias, L. F., Nothias-Esposito, M., Bouslimani, A., ... & Dorrestein, P. C. (2020). Reproducible molecular networking of untargeted mass spectrometry data using GNPS. Nature protocols, 15(6), 1954-1991.

[21] Wishart, D. S., Guo, A., Oler, E., Wang, F., Anjum, A., Peters, H., ... & Gautam, V. (2022). HMDB 5.0: the human metabolome database for 2022. Nucleic acids research, 50(D1), D622-D631.

[22] Djoumbou Feunang, Y., Eisner, R., Knox, C., Chepelev, L., Hastings, J., Owen, G., ... & Wishart, D. S. (2016). ClassyFire: automated chemical classification with a comprehensive, computable taxonomy. Journal of cheminformatics, 8(1), 61.

[23] Pluskal, T., Castillo, S., Villar-Briones, A., & Orešič, M. (2010). MZmine 2: modular framework for processing, visualizing, and analyzing mass spectrometry-based molecular profile data. BMC bioinformatics, 11(1), 395.

[24] Tsugawa, H., Cajka, T., Kind, T., Ma, Y., Higgins, B., Ikeda, K., ... & Arita, M. (2015). MS-DIAL: data-independent MS/MS deconvolution for comprehensive metabolome analysis. Nature methods, 12(6), 523-526.

[25] Tsugawa, H., Ikeda, K., Takahashi, M., Satoh, A., Mori, Y., Uchino, H., ... & Arita, M. (2020). A lipidome atlas in MS-DIAL 4. Nature biotechnology, 38(10), 1159-1163.

[26] Dührkop, K., Shen, H., Meusel, M., Rousu, J., & Böcker, S. (2015). Searching molecular structure databases with tandem mass spectra using CSI: FingerID. Proceedings of the National Academy of Sciences, 112(41), 12580-12585.

[27] Böcker, S., & Dührkop, K. (2016). Fragmentation trees reloaded. Journal of cheminformatics, 8(1), 5.

[28] Dührkop, K., Nothias, L. F., Fleischauer, M., Reher, R., Ludwig, M., Hoffmann, M. A., ... & Böcker, S. (2021). Systematic classification of unknown metabolites using high-resolution fragmentation mass spectra. Nature biotechnology, 39(4), 462-471.

[29] Ruttkies, C., Schymanski, E. L., Wolf, S., Hollender, J., & Neumann, S. (2016). MetFrag relaunched: incorporating strategies beyond in silico fragmentation. Journal of cheminformatics, 8(1), 3.
[30] Wolf, S., Schmidt, S., Müller-Hannemann, M., & Neumann, S. (2010). In silico fragmentation for computer assisted identification of metabolite mass spectra. BMC bioinformatics, 11(1), 148.
[31] Allen, F., Pon, A., Wilson, M., Greiner, R., & Wishart, D. (2014). CFM-ID: a web server for annotation, spectrum prediction and metabolite identification from tandem mass spectra. Nucleic acids research, 42(W1), W94-W99.
[32] Djoumbou-Feunang, Y., Pon, A., Karu, N., Zheng, J., Li, C., Arndt, D., ... & Wishart, D. S. (2019). CFM-ID 3.0: significantly improved ESI-MS/MS prediction and compound identification. Metabolites, 9(4), 72.
[33] Huber, F., Verhoeven, S., Meijer, C., Spreeuw, H., Castilla, E. M. V., Geng, C., ... & Spaaks, J. H. (2020). Matchms-processing and similarity evaluation of mass spectrometry data. BioRxiv, 2020-08.
[34] Huber, F., Ridder, L., Verhoeven, S., Spaaks, J. H., Diblen, F., Rogers, S., & Van Der Hooft, J. J. (2021). Spec2Vec: Improved mass spectral similarity scoring through learning of structural relationships. PLoS computational biology, 17(2), e1008724.
[35] Huber, F., van der Burg, S., van der Hooft, J. J., & Ridder, L. (2021). MS2DeepScore: a novel deep learning similarity measure to compare tandem mass spectra. Journal of cheminformatics, 13(1), 84.
[36] de Jonge, N. F., Louwen, J. J., Chekmeneva, E., Camuzeaux, S., Vermeir, F. J., Jansen, R. S., ... & van der Hooft, J. J. (2023). MS2Query: reliable and scalable MS2 mass spectra-based analogue search. Nature Communications, 14(1), 1752.
[37] de Jonge, N. F., Chekmeneva, E., Schmid, R., Joas, D., Truong, L. J., van der Hooft, J. J., & Huber, F. (2026). Cross ionization mode chemical similarity prediction between tandem mass spectra in metabolomics. Nature Communications.
[38] Padilla-González, F., Rizzo, S., Dirks, C., Bergkamp, W., Rasker, S., & Aloisi, I. (2025). Creation of an Open-Access High-Resolution Tandem Mass Spectral Library of 1000 Food Toxicants. Analytical Chemistry, 97(43), 23822-23830.
[39] Schymanski, E. L., Jeon, J., Gulde, R., Fenner, K., Ruff, M., Singer, H. P., & Hollender, J. (2014). Identifying small molecules via high resolution mass spectrometry: communicating confidence.
[40] Dührkop, K., Fleischauer, M., Ludwig, M., Aksenov, A. A., Melnik, A. V., Meusel, M., ... & Böcker, S. (2019). SIRIUS 4: a rapid tool for turning tandem mass spectra into metabolite structure information. Nature methods, 16(4), 299-302.
[41] Shannon, C. E. (1948). A mathematical theory of communication. The Bell system technical journal, 27(3), 379-423.
[42] Wang, M., Carver, J. J., Phelan, V. V., Sanchez, L. M., Garg, N., Peng, Y., ... & Bandeira, N. (2016). Sharing and community curation of mass spectrometry data with Global Natural Products Social Molecular Networking. Nature biotechnology, 34(8), 828-837.
[43] Mongia, M., Yasaka, T. M., Liu, Y., Guler, M., Lu, L., Bhagwat, A., ... & Mohimani, H. (2024). Fast mass spectrometry search and clustering of untargeted metabolomics data. Nature biotechnology, 42(11), 1672-1677.
[44] Damiani, T., Jarmusch, A. K., Aron, A. T., Petras, D., Phelan, V. V., Zhao, H. N., ... & Wang, M. (2025). A universal language for finding mass spectrometry data patterns. Nature Methods, 22(6), 1247-1254.
[45] Zuffa, S., Schmid, R., Bauermeister, A., P. Gomes, P. W., Caraballo-Rodriguez, A. M., El Abiead, Y., ... & Dorrestein, P. C. (2024). microbeMASST: a taxonomically informed mass spectrometry search tool for microbial metabolomics data. Nature microbiology, 9(2), 336-345.
[46] da Silva, R. R., Dorrestein, P. C., & Quinn, R. A. (2015). Illuminating the dark matter in metabolomics. Proceedings of the National Academy of Sciences, 112(41), 12549-12550.
[47] Blaženović, I., Kind, T., Ji, J., & Fiehn, O. (2018). Software tools and approaches for compound identification of LC-MS/MS data in metabolomics. Metabolites, 8(2), 31.