



HAL
open science

Automated fragment identification for electron ionisation mass spectrometry: application to atmospheric measurements of halocarbons

Myriam Guillevic, Aurore Guillevic, Martin K Vollmer, Paul Schlauri,
Matthias Hill, Lukas Emmenegger, Stefan Reimann

► To cite this version:

Myriam Guillevic, Aurore Guillevic, Martin K Vollmer, Paul Schlauri, Matthias Hill, et al.. Automated fragment identification for electron ionisation mass spectrometry: application to atmospheric measurements of halocarbons. 2021. hal-03176025v1

HAL Id: hal-03176025

<https://inria.hal.science/hal-03176025v1>

Preprint submitted on 22 Mar 2021 (v1), last revised 7 Dec 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PREPRINT

Automated fragment identification for electron ionisation mass spectrometry: application to atmospheric measurements of halocarbons

Myriam Guillevic^{1*}, Aurore Guillevic², Martin K. Vollmer¹, Paul Schlauri¹, Matthias Hill¹, Lukas Emmenegger¹ and Stefan Reimann¹

*Correspondence:

myriam.guillevic@empa.ch

¹Laboratory for Air Pollution
/Environmental Technology,
Empa, Swiss Federal Laboratories
for Materials Science and
Technology, Ueberlandstrasse 129,
8600 Dübendorf, Switzerland
Full list of author information is
available at the end of the article

Abstract

Background:

Non-target screening consists in searching a sample for all present substances, suspected or unknown, with very little prior knowledge about the sample. This approach has been introduced more than a decade ago in the field of water analysis, but is still very scarce for indoor and atmospheric trace gas measurements, despite the clear need for a better understanding of the atmospheric trace gas composition.

For a systematic detection of emerging trace gases in the atmosphere, a new and powerful analytical method is gas chromatography (GC) of preconcentrated samples, followed by electron ionisation, high resolution mass spectrometry (EI-HRMS). In this work, we present data analysis tools to enable automated identification of unknown compounds measured by GC-EI-HRMS.

Results:

Based on co-eluting mass/charge fragments, we developed an innovative data analysis method to reliably reconstruct the chemical formulae of the fragments, using efficient combinatorics and graph theory. The method (i) does not require the presence of the molecular ion, which is absent in ~40% of EI spectra, and (ii) permits to use all measured data while giving more weight to mass/charge ratios measured with better precision. Our method has been trained and validated on >50 halocarbons and hydrocarbons with a molar masses of 30 – 330 g mol⁻¹, measured with a mass resolution of approx. 3500. For >90% of the compounds, more than 90% of the reconstructed signal is correct. Cases of wrong identification can be attributed to the scarcity of detected fragments per compound (less than six measured mass/charge) or the lack of isotopic constrain (no rare isotopocule detected).

Conclusions: Our method enables to reconstruct most probable chemical formulae independently from spectral databases. Therefore, it demonstrates the suitability of EI-HRMS data for non-target analysis and paves the way for the identification of substances for which no EI mass spectrum is registered in databases. We illustrate the performances of our method for atmospheric trace gases and suggest that it may be well suited for many other types of samples.

Keywords: non-target screening; automated compound identification; combinatorics; machine learning; atmospheric trace gases; electron ionisation; time of flight mass spectrometry

Background

Non-target screening (NTS) is an emerging approach for analysing environmental samples, with potentially revolutionary outcomes. NTS aims to detect, identify and quantify substances that are unknown in a sample, with no or very little a priori knowledge. This approach contrasts with the more traditional target or suspect approaches, where a sample is screened only for compounds already known or suspected to be present.

So far, NTS has been developed mostly in the fields of drinking water monitoring, food analysis, forensics and metabolomics [e.g., 1–4], with human health or economic interests as the major underlying motivation. Yet, for the analysis of trace compounds in ambient or indoor air, only very limited NTS-related research has been done (e.g., the discovery of the greenhouse gas SF₅CF₃ [5, 6]), despite the need for a better understanding of the composition of the air. To look for emerging gases relevant for climate or air quality, suspect approaches are still nearly exclusively used [7–9].

NTS requires to measure properties that are specific for one given compound. In practice this is usually achieved by chromatographic time separation of the compounds. Further, the type of molecule ionisation and the mass range and mass accuracy are particularly relevant for NTS.

Originally, NTS was developed for medium to large molecules, therefore using soft ionisation such as chemical ionisation (CI), producing only a few relatively large fragments. As the molecular ion (entire molecule without one electron) is normally present and detected with CI, it is possible to reconstruct the chemical formula (i.e., the atomic assemblage, without any structural information) of the compound. To elucidate its structure, additional fragmentation and detection is required. Most (semi-)automated identification software packages were developed for CI so far [10–14].

In contrast, atmospheric trace gases consist of relatively small molecules which are best ionised by the hard electron ionisation (EI) technique. This causes a fragmentation cascade, producing many relatively small fragments; the resulting mass spectra contain valuable structural information but often lack the molecular ion [e.g., 15, Chap. 6]. Consequently, the identification of the original molecule becomes highly challenging. To circumvent this, measurements could be repeated using soft ionisation. However, such a combined analytical approach is expensive and time consuming.

A well-established approach to identify a compound based on its assemblage of masses measured by EI-MS or EI-HRMS, under the absence of the molecular ion, is to perform a mass spectrum library search. Indeed, EI ionisation has been standardised already in the 90's and produces reproducible mass spectra [15, e.g. Chap. 5]. One of the best known EI library is the NIST/EPA/NIH Mass Spectral Library, with more than 250'000 identified spectra, including approx. 140 spectra for C1 molecules [16]. However, only known and analysed compounds are present in these libraries, and identification results are therefore biased towards these compounds. Unknown emerging pollutant cannot be found by such library search. Structure-generation programs, followed by fragmentation programs, can also be used to produce candidate mass spectra otherwise absent from libraries [e.g., the MOLGEN program 17, 18, and references therein]. However, structure-generation programs require detecting

or guessing candidate molecular ions. If absent from the mass spectra, they are usually generated using other library searches. Here again, this procedure may fail to suggest candidate compounds that are absent from libraries.

Recent developments in mass spectrometry detection now provide sufficient information for broad, non-target screening approaches. In the years 2000, fast response, large coverage and high accuracy mass analyser, such as Orbitrap and time-of-flight (ToF) mass spectrometers, were introduced for water analysis, but only recently, first approaches have been made to use these powerful detectors also for organic atmospheric trace gases [19, 20]. Due to the challenge of identifying compounds, EI-HRMS are currently used as large mass-range coverage target or suspect screening instruments [e.g., 21] but only rarely as NTS instruments. For state-of-the art EI-HRMS, there is currently a huge divide between what it can deliver in terms of sample coverage, throughput and mass resolution, compared to what identification tools can provide.

In this article, we present a workflow to correctly reconstruct the chemical formula of fragments produced by the fragmentation of a precursor molecule in GC-EI-ToFMS. In addition, we develop a ranking method to identify most probable solutions and the reconstructed fragments that are most similar to the molecular ion. We evaluate our method by quantification of the correct results, on a training set of molecules and on an additional validation set. We show that reconstructing these experimental properties is possible without any library search, i.e. without favouring compounds present in databases, which pave the way for compound identification in an unbiased manner.

1 Experimental data

1.1 Training data set

To develop our methodology, we use known compounds routinely measured within the Advanced Global Atmospheric Gases Experiment (AGAGE) network [22], reported in Table 1. Most of the substances are halocarbons, i.e. molecules made of a carbon chain, with halogen atoms, and are present in the atmosphere as trace gases. Structures include saturated and unsaturated chains and the presence of rings.

Within AGAGE, the chromatographic and mass spectrometric properties are obtained by measuring diluted aliquots of a pure compound [23–26] (identification at Level 1 according to the classification for non-target analysis introduced by Schymanski *et al.* [27]). Subsequently, an unbroken chain of calibration from the prepared synthetic primary standards to measurements on our instrument ensures that the correct compounds are measured, with the correct quantification [28].

1.2 Validation data set

To validate the model after its training phase, we use a set of potentially emerging compounds, listed in Table 2. We prepared a qualitative standard containing 18 new hydrofluorocarbons (HFCs), listed under the Kigali Amendment to the Montreal Protocol [29]. The use of these substances will be progressively restricted in the coming years. Developing the capacity to check for their presence in the air, and their future molar fraction decrease, is part of supporting the application of the Kigali Amendment. The preparation of the qualitative standard is described in the Supplement. In addition, we use three hydrofluroolefins (HFOs) newly detected in air [7], which are replacing the HFCs in applications such as foam blowing and refrigeration [e.g., 30]: HFO-1234yf, HFO-1234ze(E) and HCFO-1233zd(E). We use already available standards prepared for these HFOs [7, 26].

1.3 Measurement by GC-EI-TOFMS

Our measurement system is very similar to earlier setups [25, 28]. In brief, it starts with a preconcentration trap, refrigerated at approx. -150°C using a Stirling engine, able to concentrate trace gases from up to six litres of gas (atmospheric air or reference gas mixture). Stepwise thawing of the trap eliminates the most abundant air constituents, carbon dioxide and methane, and any remaining oxygen or nitrogen, that would otherwise saturate the detector. Remaining compounds are separated by a gas chromatograph (GC), equipped with a Gaspro pre- and main column (5 m and 60 m, respectively, 0.32 mm inner diameter, Agilent), ionised and detected by a time-of-flight detector (H-TOF, Tofwerk AG, Thun, Switzerland). The detector is set to measure fragments with masses from 24 m/z to 300 m/z. Masses below 24 m/z are prevented from hitting the detector, to avoid saturation by potential water contamination. The mass resolution is approximately 3000 below m/z of 50 and up to 4000 above m/z of 100. The raw intensity data at each time-of-flight and each time bin are saved in a file of format `hdf5` [31], which constitutes the used raw data. The total analysis time for one sample is 70 min, with 40 min of preconcentration and stepwise thawing, followed by 30 min of gas chromatography and detection by TOFMS.

Intensity data, defined as the number of ions that hit the detector at a certain time, are recorded along two axes, the time-of-flight axis (later converted to a mass

axis) and the retention time (RT) axis. While the signal along the RT axis reflects the separation of molecules by the GC, the signal along the mass axis reflects the fragmentation pattern of the molecules measured by EI-MS. One fragmented, detected molecule can be visualised as a mountain ridge, producing a variety of mass peaks with various intensities, all aligned perpendicular to the time axis. First, peaks are detected and fitted along the mass axis, and afterwards along the RT axis.

Along the time-of-flight axis, at each time bin, peaks are fitted using a pseudo-Voigt function (Supplement). The obtained time-of-flight centres of the peaks are then converted to masses, using the mass calibration function (Supplement). We assume that all masses have been ionised just once. This produces a set of 20 to 30 centres of mass values with associated RT and intensity, from which the mean and standard deviation are computed, weighted by intensity. For each detected peak, the mass uncertainty (u_{mass}) is the Euclidean distance of the calibration uncertainty ($u_{\text{calibration}}$, see Supplement) and the measured standard deviation ($u_{\text{measurement}}$):

$$u_{\text{mass}} = 2.5\sqrt{u_{\text{calibration}}^2 + u_{\text{measurement}}^2} \quad (1)$$

This u_{mass} is computed using a coverage factor of 2.5 to constrain the range of possible masses for the knapsack algorithm described below. This corresponds approximately to a 98.5% confidence interval.

Along the RT axis, data are saved with a frequency of six points per second (6 Hz). Usually, chromatography peaks last for a minimum of four seconds, producing 20 to 30 points per peak in the RT domain. The observed peak shape along the time axis is typical for gas chromatography and is fitted using the equation proposed by Pap et al. [32], that in our case fits well the observed tailing. Finally, co-eluting mass peaks are grouped together.

Routine quality control of instrumental performance includes measuring blanks to check for potential contaminants coming from the measurement setup itself, drifts in retention time due to column ageing or water contamination, and stability of intensity ratios of mass fragments belonging to the same compound.

2 Method for automated fragment formula identification

2.1 Method overview

The output after peak fitting and mass calibration is a dataset of mass/charge ratio (m/z), each with intensity (in V) and uncertainty (in ppm), at a precise retention time (in seconds). Co-eluting peaks may correspond to chemical fragments of a unique molecule, or a small number of distinct molecules. They are therefore grouped into one time slot of approx. 2 s width to be treated together by the identification algorithm. We consider each time slot separately.

The overview of our method is depicted in Fig. 1. The general approach is to consider separately each group of co-eluting fragments, and to reconstruct the chemical formula of each fragment based on two types of information:

- from the experimental data produced by GC-EI-ToF analysis, we use the measured mass and measured signal intensity of each peak. For the measured masses, the uncertainty (Eq. (1)) is computed following metrology principles (Fig. 1, yellow box **Input**).

- we combine these experimental data with chemical information that is universal, i.e. true for any given molecule: exact mass and valence of chemical elements, known environmental stable isotopic abundances (Fig. 1, two mauve boxes).

In practice, the identification method combines algorithms for two purposes: (i) algorithms that enumerate solutions in an exhaustive way, according to given constraints (Fig. 1, steps 1 and 3); (ii) algorithms that eliminate unlikely solutions, based on other constraints (Fig. 1, steps 2, 7, 8 and 9).

The developed workflow is as follow (steps are numbered as in Fig. 1):

- Step 1: To start, possible atom assemblages matching the measured masses, within uncertainty, are exhaustively generated. This step usually produces a large number of possible chemical formulae;
- Step 2: all generated formulae are organised in a pseudo-fragmentation graph. This step relies on the specificity of EI-MS, producing many various fragments from the same precursor ion molecule. The irrelevant or unlikely formulae are discarded;
- Step 3: isotopocules (i.e. molecules having the same type and number of atoms, but where at least one atom is a different isotope) are generated;
- Step 4: for each set of isotopocules, the isotopic pattern of fragments (theoretical intensity profile) along the mass axis is generated. The optimum contribution to the measured profile, of each set separately, is computed;
- Step 5: using a specifically designed likelihood estimator, all candidate chemical formulae are set out in order of preference, number one being the most likely, and the last one being the least likely;
- Step 6: This ordering enables the selection of the most likely candidate(s);
- Step 7: the intensity of each isotopocule set is optimised by comparison to the measured mass profile, using a machine learning technique. This is by far the computationally most costly step. Isotopocule series characterised by a total intensity below a given threshold, usually, the instrumental limit of detection (LOD), are eliminated.
- Step 8: The pseudo-fragmentation tree is updated, and the optimisation procedure resumes at [Step 5](#), until a predefined fraction of the measured signal is reproduced.
- Step 9: Lastly, all remaining, not optimised candidate chemical formulae are deleted. The remaining candidates constitute the final list of most likely correct chemical formulae. Each measured mass may have zero, one or several assigned chemical formulae.

Each step of the method is explained in details hereafter. In the Supplement, Sect. 6, we give a numerical example with the mass spectrum that will turn out to correspond to CCl_4 .

2.2 [Step 1](#): Generating fragment formulae: the knapsack algorithm

The aim of the *knapsack step* is to recover all chemical formulae that could correspond to each detected fragment, given its mass and uncertainty, and excluding all other chemical formulae that would not fulfil the criteria of measured mass and uncertainty. This knapsack step produces the correct chemical formulae, usually along with many other incorrect formulae. The aim of subsequent steps will be to eliminate the incorrect formulae using additional constraints.

We use combinatorics to generate the chemical formulae of candidate fragments, for each mass detected in a spectrum. In particular, we develop a variant of the knapsack algorithm [33, 34], dedicated to our setting, which will be described below.

2.2.1 The knapsack algorithm in combinatorics.

Combinatorics are mathematical algorithms of fast and exhaustive enumeration, and the knapsack problem is a well-known topic in this area (see e.g. the handbook on algorithms [35]). The problem is usually stated as follows: given a knapsack of maximum available capacity (e.g. mass), and a set of items each of a specific capacity, find subset(s) of items that can fit into the knapsack, while optimising some other quantity (usually maximizing the total price of the items). In our setting, the knapsack is a fragment of given mass (within the uncertainty range), and the items are atoms of exact given mass. We are interested in enumerating all possible combinations of atoms so that the sum of their masses fits the measured fragment mass within the uncertainty margin. An unbounded number of each atom is available i.e. each atom type can be used multiple times, this is also known as the *unbounded* knapsack problem. Contrary to the classical problem in combinatorics, we do not optimise another parameter. Instead, we are interested in listing all possibilities. In this work, we design an algorithm for a fast and exhaustive enumeration of all the solutions to the knapsack problem.

The inputs of our dedicated knapsack algorithm are the measured masses of the fragments with uncertainties, and a list of masses of elements that are expected to form the fragments. Because the exhaustive list of solutions grows exponentially with the number of elements, we will introduce different techniques to avoid as early as possible enumerating wrong chemical formulae, while still being exhaustive.

2.2.2 Targeted mass with uncertainty.

This section describes the basic algorithm to enumerate all the possibilities. The target mass, which is one measured mass, is denoted by m ; the set of item masses, which are the exact masses of chemical elements, is made of I distinct positive values m_i , labelled m_0 to m_{I-1} , sorted in increasing order, that is, $m_i < m_{i+1}$ for all i . We consider IUPAC periodic table values [36]. We do not consider uncertainties of the atomic masses, which are negligible compared to the ToF analytical mass uncertainties. A solution of a knapsack problem is encoded as a vector of positive integers $[a_0, a_1, \dots, a_{I-1}]$ where a_i is the number of items of mass m_i ; it is 0 if the item i is not in the solution. Algorithms 1, 2 (in pseudo-code) describe the basic recursive enumeration. An iterative (non-recursive) function is also possible but we implemented a recursive function.

2.2.3 Only the most abundant isotope of each element used as input.

We consider 9 elements with their stable isotopes (if any), making a list of 19 different exact masses [1] that can be combined to form a molecule (we omit the elements that are rarely found in volatile atmospheric trace gases, such as metals). The electron ionisation fragmentation produces isotopocule fragments. For example for the molecule CCl_4 , we may observe CCl_4 made of only abundant isotopes (^{12}C

[1]H, ^2H , C, ^{13}C , N, ^{15}N , O, ^{17}O , ^{18}O , F, S, ^{33}S , ^{34}S , Cl, ^{36}S , ^{37}Cl , Br, ^{81}Br , I.

Algorithm 1: `unbounded_knapsack_rec` header function, Exhaustive enumeration of solutions to an unbounded knapsack problem

Input: positive target mass bounds m_{\min}, m_{\max} , ordered item mass list $[m_0, m_1, \dots, m_{I-1}]$ where $0 < m_i < m_{i+1}$

Output: set S of solutions $[a_0, a_1, \dots, a_{I-1}]$ such that $m_{\min} \leq \sum_{i=0}^{I-1} a_i m_i \leq m_{\max}$, $a_i \geq 0$

$S \leftarrow \{\}$ (empty set)

$\mathbf{a} = [0, 0, \dots, 0]$ (zero vector of length I)

$S \leftarrow \text{unbounded_knapsack_rec_aux}(m_{\min}, m_{\max}, [m_i]_{0 \leq i < I}, \mathbf{a}, I - 1, S)$

return S

Algorithm 2: `unbounded_knapsack_rec_aux` Recursive function

Input: positive target mass bounds m_{\min}, m_{\max} , ordered item mass list $[m_0, m_1, \dots, m_{I-1}]$ where $0 < m_i < m_{i+1}$, vector $\mathbf{a} = [a_0, \dots, a_{I-1}]$, index i ($0 \leq i < I$), partial solution set S

Output: updated set S of solutions $[a_0, a_1, \dots, a_{I-1}]$ such that $\sum_{i=0}^{I-1} a_i m_i = m \in [m_{\min}, m_{\max}]$, $a_i \geq 0$

if $i < 0$ **or** $m_{\max} < 0$ **or** ($m_{\min} > 0$ **and** $m_{\max} < m_0$) **then**
 return S

if $i = 0$ **then**
 $a_0^{\max} \leftarrow \lfloor m_{\max}/m_0 \rfloor$
 $a_0^{\min} \leftarrow \max(0, \lceil m_{\min}/m_0 \rceil)$
 for a_0 **from** a_0^{\max} **downto** a_0^{\min} **by** -1 **do**
 $\mathbf{a}[0] \leftarrow a_0$ (we have $m_{\min} \leq a_0 m_0 \leq m_{\max}$)
 Append \mathbf{a} to S
 return S

if $i > 0$ **then**
 $a_i^{\max} \leftarrow \lfloor m_{\max}/m_i \rfloor$
 for a_i **from** a_i^{\max} **downto** 0 **by** -1 **do**
 $\mathbf{a}[i] \leftarrow a_i$
 $S \leftarrow \text{unbounded_knapsack_rec_aux}(m_{\min} - a_i m_i, m_{\max} - a_i m_i, [m_i]_{0 \leq i < I}, \mathbf{a}, i - 1, S)$
 return S

or C in short notation, ^{35}Cl or Cl in short notation), and isotopocules containing ^{13}C and ^{37}Cl (see the complete list of isotopocules provided in Fig. 3 and in the Supplement, Table 8).

To reduce the enumeration of the knapsack, the input is limited to the mass of the most abundant isotope of each atom (e.g. C of mass $12.000000 \text{ g mol}^{-1}$ for carbon, Cl of mass $34.96885271 \text{ g mol}^{-1}$ for chlorine), making a list of 9 exact masses to be used for the enumeration, instead of 19. For relatively small molecules, the fragment made of only abundant atoms has usually the highest intensity (or second highest, Fig. 3), hence producing a much smaller mass uncertainty than for the other isotopocules. The target mass range is thinner, reducing the knapsack enumeration. By doing so, we obtain possible solutions with the knapsack only for some of the most abundant fragments. Once a fragment of abundant atoms is generated by the knapsack, we later enumerate all its isotopocules containing rare isotopes, in Step 3 (Fig. 1). This is further explained in Section 2.5.

2.2.4 Optimisation of the knapsack enumeration: double bond equivalent (DBE) criterion with meet-in-the-middle optimisation algorithm.

Not all sets of atoms form a valid chemical formula. Indeed, each atom allows a maximum number of chemical bonds with other atoms, according to its valence. This can be expressed using the *double bond equivalent* (DBE), or sum of number of rings and double bonds in a given chemical formula. The DBE is computed with [§6.4.4

Eq. (6.9) 15]

$$\text{DBE} = 1 + 0.5 \times \sum_i^{i_{\max}} N_i (V_i - 2) \quad (2)$$

where V_i is the valence of element i and N_i is the number of atoms of element i in the chemical formula. For EI-MS, since we expect no cluster formation in the ionisation source, the minimum valence for a chemical formula is 0. We do not set any maximum valence value. For the sulphur element, where multiple valences are possible, we chose the maximum value (6), according to one of the golden rules for identification [37].

Of all chemical formulae matching the considered mass domain, only a fraction are chemically valid. This means that to reduce the enumeration time of the knapsack, one strategy is to avoid enumerating chemical formulae with a negative DBE value. We explain hereafter how we implement this.

In the early 80's, cryptologists^[2] formulated a meet-in-the-middle strategy to speed-up the enumeration of all solutions of a knapsack problem [33]. The key-ingredient is to partition the candidate items in two sets. Applying this strategy to our topic, one enumerates all possibilities made of items of the first set and whose mass is smaller or equal to the target mass. The possibilities are ordered by increasing mass. Meanwhile, one does the same with the items of the second set. The two sets are chosen so that the respective running-time of the two enumerations are balanced, in order to minimize the total running-time. We end up with two lists of masses in increasing order, of value between 0 and the target mass. Then reading onward the first list and downwards the second list, one looks for pairs of partial solutions, one from each list, so that the paired mass matches the target mass. A numerical example is given in the Supplement in Section 6.1.

We adapt this strategy to speed up the solution of our problem. We partition the input atoms in two sets: the set of multivalent atoms (C, N, O, S) and the set of monovalent atoms (H, F, Cl, Br, I). First, all solutions made of multivalent atoms, and smaller or equal to the target mass, are generated. To a generated multivalent-atom-solution, a maximum number of monovalent atoms, twice the DBE value (cf. Eq. (2)), can be added and still form a valid chemical formula. In this way, the partial solutions made of multivalent atoms give us an upper bound on the number of monovalent atoms that can complete the fragment, reducing considerably the enumeration of partial possibilities with monovalent atoms. In particular, it gives an upper bound on the number of hydrogen H. A numerical example is given in the Supplement in Section 6.2.

The list made of multivalent atoms is precomputed for the heaviest mass first, and can be re-used for the smaller masses. We also implemented a way to re-use the list of monovalent atoms precomputed for the heaviest target mass, when processing the lighter target masses. Our strategy turned out to be fast enough for the considered mass ranges (see runtime in Sect. 3.5) and we did not investigate

^[2]A variant of the knapsack problem was used to build cryptosystems to securely hide secrets in the 70's. It was later broken with the LLL algorithm. We leave to future work the application of the LLL algorithm to our setting.

further optimisations. Dührkop et al. [38] have a very different approach well-suited for mass ranges of around 1000 m/z.

After **Step 1** (Fig. 1), for each measured mass, all chemical formulae that are in agreement with the measured mass within its uncertainty, made of abundant isotopes, and having a positive DBE value, are generated. At this stage, the fragment formulae are not uniquely identified by the knapsack: for each measured mass there are either too many possibilities, or none (usually because the fragment may contain non-abundant isotopes).

2.3 **Step 2:** Organisation of the results in a pseudo-fragmentation graph

The aim of **Step 2** is to organise all chemical formulae generated in **Step 1** according to a specific order, to help identify and delete unlikely chemical formulae.

With EI-MS, a fragmentation cascade happens due to the high ionisation energy, i.e. several fragmentation steps one after the other [e.g., 15]. A fragmentation step produces an ionised fragment (detected) and a neutral (not detected). Each detected fragment may result from one or several fragmentation steps. As the EI fragmentation takes place under vacuum with pressure usually below 10^{-5} bar, we do not expect to see clusters originating from agglomeration of (fragments of) the molecular ion with other chemical species. On the contrary, all detected fragments are pieces of the original molecule. If all fragmentation steps are known, one can organise the fragments in an acyclic directed graph (Fig. 2). The nodes are the fragments. One edge is one fragmentation step. This forms a *fragmentation graph*. Potentially, several fragmentation pathways in the EI source may produce the same end fragment(s). But thanks to directions, the graph is acyclic.

We organise all the candidate formulae from the knapsack in a directed graph (with the class `DiGraph` provided in the Python package `Networkx`, [39]). Each node on the graph is a candidate fragment, with associated attributes, such as its chemical formula, its exact mass, the associated measured mass(es), and the list of rare isotopocules that will be generated at the next step. An edge is set from a node to another if the chemical formula of the first fragment admits the chemical formula of the second one as subformula (e.g. CCl_3 has subfragment CCl , see Fig. 2 presenting the directed graph obtained with all knapsack solutions of CCl_4). This mimics the possible fragmentation pathways. In other words, we define a *partial* ordering of the fragments (it is not *total* because, for example, there is no relation between candidate fragments CCl_3 and CSBr , cf. Fig. 2). The *maximal fragments*^[3] have no ancestor but may have children (Fig. 2, nodes in orange and yellow), they are the maximal elements of the ensemble of fragments. If the molecular ion is present, it is one of the maximal fragments. As with EI, the molecular ion is often absent (as for 15 compounds of the training set, see Table 1 in the Supplement), several maximal fragments are allowed. Also, to account for potential co-elution of different molecules, several connected components are allowed. This algorithm does not use any structural information, only the candidate chemical formulae, producing only a pseudo-fragmentation graph, not a chemically realistic one, contrary to previous algorithms [40]. We do not use a list of expected neutral losses (as in e.g., [11])

^[3]in the usual mathematical meaning, e.g. https://en.wikipedia.org/wiki/Partially_ordered_set#Extrema

due to the high heterogeneity of our chosen molecules. Therefore, it is likely that some edges are actually not structurally possible, but this is not relevant at this stage. Optimisation strategies for an efficient construction of this directed graph are reported in the Supplement (Sect. 5.1).

From this pseudo-fragmentation graph, we would like to eliminate the isolated nodes, or nodes not being connected to any other node (singletons), with neither ancestors nor children (Fig. 2, nodes in yellow). They may correspond to (i) impurities produced for example by GC bleed or (ii) solutions from the knapsack that are unlikely, in particular with an atom absent from all other nodes. But we need to account for very small molecules such as CFC-11 (CFCl_3) or CFC-13 (CF_3Cl) that produce a very limited number of different fragments, without the molecular ion: if one measured mass has only singletons as candidate fragments, we do not eliminate them.

At the end of [Step 2](#), usually a few unlikely knapsack solutions being singletons have been eliminated.

2.4 [Step 3](#): Enumeration of chemical formulae containing rare isotopes

The knapsack algorithm produces candidate chemical formulae made of abundant isotopes only. But all isotopocules of a fragment are expected to be present in a given time slot (the very rare ones which are below the detection limit of the mass spectrometer may not be detected). Therefore, for each candidate chemical formula, we now generate a set of all isotopocules including their relative intensities based on their natural isotopic abundances [[36](#)]).

Hereafter, we name *knapsack formula* a chemical formula from the knapsack, and *rare-isotope formula* a chemical formula with at least one rare isotope, even if this isotopocule is expected to be more abundant than the *abundant chemical formula*. For example, CCl_4 is called knapsack formula, while $\text{CCl}_3[^{37}\text{Cl}]$ is called rare-isotope formula.

For each possible knapsack formula, we generate the list of all possible isotopocules. Again, this is an enumeration process using combinatorics. If the chemical formula is made of atoms that are monoisotopic, the list contains the knapsack formula only, whose intensity is one, that is, 100%. Otherwise, for each rare-isotope formula, we compute its exact mass and expected intensity relative to the knapsack formula [[36](#), [41](#)]. The isotopocules of the knapsack formula CCl_4 with their relative intensities are shown in Fig. 3 and the corresponding numerical values can be found in Table 8 of the Supplement.

Optimisation strategies to speed up the enumeration are reported in the Supplement (Sect. 5.3). In particular, isotopocules with intensity expected below the instrumental limit of detection are not enumerated.

At the end of [Step 3](#), knapsack-generated chemical formulae, containing only abundant isotopes, are organised as nodes in a pseudo-fragmentation graph. Each knapsack formula of a node is complemented by its rare-isotope formula(e) if the latter is above the instrumental limit of detection.

2.5 [Step 4](#): Computing the optimum contribution for each isotopocule set individually

We now consider the measured mass intensity profile. Potentially, any candidate chemical formula may contribute to the measured intensity profile along the mass

axis. First, one generates the theoretical mass profile for each node, i.e. for each knapsack formula together with its rare-isotope formulae. Then, one optimises a certain contribution for each node taken individually, to match the measured mass profile.

2.5.1 Computing a profile of intensity vs mass for a given set of isotopocules

For each set of isotopocules belonging to the same node, we generate an expected mass profile. A measured intensity profile is not continuous, it is a discrete set of coordinates (m_a, I_{m_a}) where m_a is a mass abscissa, and I_{m_a} is the intensity for this mass. We consider that a knapsack fragment and its associated rare-isotope fragments have an expected intensity profile made of the sum of contributions of all isotopocules along the mass range. Each mass peak is generated as a pseudo-Voigt profile, with a prescribed peak width as obtained from the mass calibration (Supplement) and a mass resolution of about 8 ppm, which is sufficient given our instrumental mass resolution. We obtain an expected discrete mass profile (a set of coordinates) for the whole isotopocule set of the node, of the form $\{(m_a, \tilde{I}_{m_a}) : m_a \in \text{mass abscissa}\}$.

2.5.2 Computing the contribution of a given set of isotopocules

At this point, over the mass domain covered by a given candidate set of isotopocules i , we look for a non-negative scaling factor k_i , such that the measured signal $s_{\text{measured}} = \{(m_a, I_{m_a}) : m_a \in \text{mass abscissa}\}$ best fits the theoretical profile of the set. This can be seen as an optimisation problem, where the difference between measured and generated profile is minimised:

$$s_{\text{measured}} - k_i \cdot \tilde{I}_i = \sum_{m_a \in \text{mass abscissa}} (I_{m_a} - k_i \cdot \tilde{I}_{i,m_a}) = 0 \quad (3)$$

where s_{measured} and \tilde{I}_i can be seen as vectors. With only one k_i value to optimise, Eq. (3) can be simplified as the average value of the measured profile divided by the computed isotopocule profile:

$$k_i = \frac{1}{\#\{m_a \in \text{mass abscissa}\}} \sum_{m_a \in \text{mass abscissa}} \frac{I_{m_a}}{\tilde{I}_{i,m_a}} \quad (4)$$

After optimisation of the factor k_i , if the entire profile falls below the LOD, the candidate node is removed from the graph of solutions.

This step can be seen as computing the maximum contribution for a given node (as if no other nodes were present to contribute to the signal). The computed maximum contribution, that we denote k_i^{max} , will be used as starting value for the first computation of likelihood in [Step 5](#).

2.6 Step 5: Ranking of candidates according to a likelihood estimator

To help us select the most probably correct formulae among all candidate formulae, we define an artificial likelihood estimator based on two quantifications, as explained hereafter. The estimator takes as input a knapsack fragment, together with its set of rare-isotope fragments.

We decide to capture in the likelihood estimator the other knapsack fragments that are sub-fragments, and their isotopocules. For example, for CCl_2 as chosen knapsack formula, we would consider its five rare-isotope formulae (e.g. $^{13}\text{C}]\text{Cl}_2$), the subformulae CCl and Cl and their own set of isotopocules (C of m/z 12 is filtered out in our ToF MS). At [Step 4](#), we estimated k_i^{\max} , the maximum value that k_i could take. We now estimate the maximum proportion g of signal that a candidate fragment n , its isotopocules, and all its sub-fragments i could explain. These considerations lead to a first possible estimator:

$$g(n) = \frac{\sum_{i \text{ sub-fragment of } n} k_i^{\max} \sum_{j \text{ isotopocule of } i} p_{i,j}}{\sum_{m_{\text{spec}} \in \text{mass spectrum}} I_{m_{\text{spec}}}} \quad (5)$$

where $p_{i,j}$ is the theoretical abundance of the given isotopocule j for the given fragment i computed with Eq. 7 in the Supplement, m_{spec} ranges over the mass spectrum, $I_{m_{\text{spec}}}$ is the intensity at that mass ($I_{m_{\text{spec}}}$ is computed as a discrete integral, this is the area of the peak of the measured signal around m_{spec}). Numeric values of $g(n)$ for knapsack fragments of CCl_4 are given in [Table 3](#), fourth column. In practice, we have observed a misbehaviour of the estimator $g(n)$ from Eq. (5): knapsack formulae that contain many atomic elements have many more sub-fragments, of various masses, hence a higher probability to match a larger fraction of the signal. Eq. (5) therefore gives advantage to “complicated” formulae. This effect can be seen for knapsack fragments of CFC-11 ([Table 3](#), col. 4): fragment HCFCl (wrong) gets a higher score than CFCl (correct) using Eq. (5).

To correct for this effect, we multiply Eq. (5) by the number of found sub-formulae, divided by the expected maximum number of sub-formulae. In practice, this maximum number of sub-formulae is computed with the knapsack algorithm, using the minimum detectable mass as lower bound (in our case, $m/z = 23$, as all smaller masses are filtered out in our ToF detector). As the number of knapsack solutions increases with i) the total number of atoms and ii) the number of elements present in the fragment, using this number as denominator will favour solutions constituted of a limited number of elements. The advantage of this technique is to favour “simple” solutions, without setting any parameter to limit the number of elements to use. Equation (5) is therefore completed as:

$$g(n) = \frac{\sum_{i \text{ sub-fragment of } n} k_i^{\max} \sum_{j \text{ isotopocule of } i} p_{i,j}}{\sum_{m_{\text{spec}} \in \text{mass spectrum}} I_{m_{\text{spec}}}} \frac{\#\{\text{sub-fragments } i \text{ of } n\}}{\#\{\text{all theoretical sub-fragments of } n, \text{ mass } \geq 23m/z\}} \quad (6)$$

where $\#\{\text{sub-fragments } i \text{ of } n\}$ is the number of existing subfragments in the directed graph, for fragment n .

All knapsack fragments are set out in order by decreasing likelihood value. Looking again at the same example in [Table 3](#), using Eq. (6) now fragment CFCl (correct) gets a higher likelihood score than HCFCl (col. 5), and is therefore ranked better (col. 6). We underline that this defined likelihood estimator has no chemical signification. Its aim is only to highlight the simplest knapsack solutions that explain the maximum proportion of the measured signal. It is therefore only a practical tool to speed up the identification process.

2.7 Step 5 to Step 8: Iterating loop to compute the optimum contribution of multiple sets of candidates together

Overall, the measured signal profile $s_{\text{measured}} = \{(m_a, I_{m_a}) : m_a \in \text{mass abscissa}\}$ should match a linear combination of the expected profiles of all correct candidate sets, an approach already found in e.g. [42]. Formally, this approach allows several fragments to form together the signal of one measured peak, which is realistic given our mass resolution. Equation (3) is therefore completed as:

$$s_{\text{measured}} - \sum k_i \cdot \tilde{I}_i = \sum_{m_a \in \text{mass abscissa}} \left(I_{m_a} - \sum_{\text{profile } i} k_i \cdot \tilde{I}_{i, m_a} \right) = 0 \quad (7)$$

This equation cannot be simplified. Instead, the optimisation of multiple k_i is handled by a machine learning technique, using the Python package `lmfit`, based on the Levenberg-Marquardt algorithm [for details, see 43]. This optimisation of multiple k_i together is by far the most expensive computation step in the entire method. The `lmfit` package is most efficient when only a limited number of parameters k_i are optimised. In particular, it is outside the computing power of a regular desktop machine to optimize all the profiles of the candidate formulae together. We implement three techniques to lower the running time of `lmfit`.

First, the measured profile is itself not perfect but affected by measurement noise. To account for this, we do not aim at reconstructing 100% of the signal, but only a significant portion of it. Since our experimental precision is in the order of 1 % to 5%, we set the threshold at 95% of the signal: when 95% of the “area” below the signal is reconstructed, the optimisation is stopped. The fraction of explained signal f is calculated as follow, iterating over the i selected candidate fragments, their respective set of isotopocules, and the mass spectrum:

$$f = \frac{\sum_{\text{selected candidate } i} k_i \sum_{\text{isotopocule } j \text{ of } i} p_{i,j}}{\sum_{m_{\text{spec}} \in \text{mass spectrum}} I_{m_{\text{spec}}}} \quad (8)$$

where k_i are the linear factors optimised by the `lmfit` package. We aim at reaching $f \geq 0.95$.

Second, a further reduction of computation time is achieved by splitting the mass domain: all candidates are grouped into smaller, non-overlapping mass domains, where the optimisation is run separately. Indeed, optimising a small number of k_i multiple times is more efficient than optimising a large number of k_i just at once.

Third, we observe that at this stage, many wrong solutions generated by the knapsack are still present, while usually, only a limited number of chemical formulae are really present (see later discussion in Sect. 3.5 and also Fig. 8). We therefore adopt the following greedy-like strategy^[4], in order to reduce the number of fitted isotopic profiles. The most likely solutions are processed first, until the reconstructed signal reaches 95% of the measured signal. This way, unlikely solutions left after 95% of the signal has been reconstructed are not considered at all. This approach

^[4]https://en.wikipedia.org/wiki/Greedy_algorithm

requires all candidate fragments to be ordered according to a well chosen likelihood estimator, as done in [Step 5](#).

In practice, from the list of ranked knapsack formulae (or nodes on the graph), we take the one ranked first, together with all its sub-fragments (or children of the node) and all associated isotopocules, optimise the contributions of these selected nodes ([Step 7](#)), and then eliminate any node below the LOD, as well as any node becoming a singleton ([Step 8](#)). Since the k_i have updated values, the ranking is updated (back to [Step 5](#)). If less than 95% of the signal is reconstructed, the next most likely node is added to the selected nodes ([Step 6](#)). This iterating procedure is depicted in [Fig. 1](#), green boxes.

At the end of this iterating procedure, nodes that have not been selected to be optimised are deleted from the pseudo-fragmentation graph, and any new singletons are deleted as well ([Step 9](#)). For each remaining node, its factor k_i is used to compute the final contribution of this node to each measured mass peak.

3 Results and discussion

3.1 Validation data from standard measurements

Proving that our methodology reconstructs the correct chemical formula(e) with the correctly assigned intensity, for each measured mass, requires an independent source. However the NIST database for example contains only spectra of unit mass, without information about fragment formulae. We therefore adopt the following pragmatic approach, already used in e.g. [\[44\]](#): a fragment formula is considered to be correct if it is a sub-formula of the molecular formula. Fragments resulting from re-arrangements, which do happen with EI ionisation, are thus considered as correct solutions.

3.2 Estimators of method performance

A first quantitative measurement of method performance is how much of the measured signal $I_{m_{\text{spec}}}$ is reconstructed, written $f_{I,\text{total}}$, irrespective of whether the signal corresponds to correct or incorrect fragments:

$$f_{I,\text{total}} = \frac{\sum_{m_{\text{spec}} \in \text{mass spectrum}} I_{\text{correct fragment}} + I_{\text{incorrect fragment}}}{\sum_{m_{\text{spec}} \in \text{mass spectrum}} I_{m_{\text{spec}}}} \quad (9)$$

Reconstruction of fragment formulae produces a qualitative information. To quantify method performance, we construct two metrics based on recommendations for qualitative measurements [\[45, Annex D\]](#).

We define the ratio of correctly identified fragment formulae f_n as the number of correctly identified fragment formulae divided by the total number of reconstructed fragment formulae, per compound:

$$f_{n,\text{correct}} = \frac{n_{\text{correct fragment}}}{n_{\text{correct fragment}} + n_{\text{incorrect fragment}}} \quad (10)$$

Then, we define the ratio of correctly assigned signal $f_{I,\text{correct}}$ as the sum of intensities of correctly assigned fragments divided by the total intensity of all reconstructed

fragments:

$$f_{I,\text{correct}} = \frac{\sum I_{\text{correct fragment}}}{\sum I_{\text{correct fragment}} + \sum I_{\text{incorrect fragment}}} \quad (11)$$

For comparison, we also compute per training compound the fraction of correct fragments and signal on the "top 10" results, which are the list of maximum 10 fragments with maximum likelihood.

3.3 Performances of the method on the training set

An example of the output produced by the identification method is given for CFC-11 in Table 3 and the corresponding mass spectra is plotted in Fig. 4. What is delivered to the user is a list of the chemical formulae of the generated and non-eliminated fragments, their exact mass, assigned signal, likelihood value according to Eq. (6) and ranking. The method also informs if a fragment is a maximal fragment. The correctly reconstructed chemical formulae of the maximal fragments are closest to the molecular ion. In fact, for CFC-11 the maximal fragment CFCl_3 with ranking of 1 (maximum likelihood) is the molecular ion (Table 3). This example demonstrates the usefulness of our likelihood estimation to detect the fragments closest to, or being, the molecular ion. Note that if the molecular ion is not detected, our method does not reconstruct it; we keep this development for future work.

3.3.1 Fraction of reconstructed signal

For all but two compounds in the training set, the fraction of reconstructed signal is above 0.95. For CF_4 and C_2H_6 , it is 0.06 and 0.88, respectively. In these two cases, the measured mass and its uncertainty envelope did not contain the true mass, causing the knapsack to fail in generating the correct chemical formulae. Given that we multiply the mass uncertainty with a coverage factor of 2.5 (Eq. (1)), corresponding to 98.5% of the expected mass interval, it can be expected that in a few cases, the considered mass domain does not contain the true answer. We observe that no wrong fragment fills this gap, but the signal is rather not reconstituted. It is therefore easier for the user to identify such extreme cases, and e.g. run the identification process again using a larger coverage factor.

For CF_4 , correct chemical formulae were reconstructed for the three measured masses when using a coverage factor of 6.0. For C_2H_6 , using $k = 2.5$, no solution is produced for the measured fragment at 29.037805 m/z. Using a coverage factor k of 3.0 instead allows the knapsack to produce the correct fragment formula, C_2H_5 .

3.3.2 Compounds with very few measured masses

As a consequence of their simple molecule structure or a very low abundance in the sample, five compounds had fewer than 6 detectable masses (NF_3 , CF_4 , CH_3I , SO_2F_2 , SF_5CF_3). We observed that in such cases, our identification algorithm does not have enough constraints to suggest correct results. We therefore developed the following strategy: when the number of measured masses is less than 6, maximal fragments are treated separately through the iterative Step 5 to Step 8 (green boxes on Fig. 1), so that chemical formulae belonging to different maximal fragments are not optimised together. The list of kept maximal fragments is then returned as

result, ordered by likelihood. The program returns a message warning that multiple maximal fragments are possible, and suggests the user to choose the one considered most likely.

For the sake of including these compounds with all other results in the following figures, the maximal fragment ranked first is kept, and all others are eliminated (assuming no co-elution). In all cases except for SO_2F_2 , where only two masses were measured, the first ranked maximal fragment was correct.

3.3.3 Fraction of correct reconstructed fragments and signal

Fig. 5 displays the histograms of the fractions of correct fragment formulae and of correct reconstructed signal, for all fragments (left) and for the top-10 fragments for each compound (right). Fig. 5 shows better fraction of correct results when taking the fraction of correct signal into account compared to the fraction of correct chemical formulae: wrongly identified fragments tend to have a smaller abundance, mostly due to higher mass uncertainty or lack of companion peak that would provide an isotopic constrain. For only 44% of the compounds, 90% or more of the chemical formulae are correct. However, for more than 90% of the compounds, the signal from correct fragments constitutes at least 90% of the reconstructed signal. This underlines that the proportion of signal assigned to each chemical formula carries information about how likely this chemical formula is correct.

The ability of the method to produce correct chemical formulae further improves when taking into account the top-10 results only, i.e. the 10 chemical formulae ranked as most likely according to the likelihood estimation, and their associated signal. For these top-10 fragments, the chemical formulae are at least 90% correct for 58% of the compounds; the proportion of compounds for which at least 90% of the signal is correct stays unchanged, at 90%. Indeed, most of the time, wrong fragments are anyway assigned a small portion of the signal.

For the training set, only two compounds were poorly identified, CF_4 and SO_2F_2 , for reasons discussed previously.

3.3.4 Information from the likelihood estimation and ranking results

As we have seen, the identification algorithm does not eliminate all wrong chemical formulae. We now study more in detail the likelihood value and ranking associated to each reconstructed fragment to see if these values can better inform us if a fragment is correct or wrong.

Fig. 6 (left histogram) presents the distribution of likelihood values for correct fragments (in blue) and for incorrect fragments (in red), and the same for maximal fragments only (right histogram). From these distributions, we observe that likelihood values above 20 indicate that the fragment is correct by 95% ($n = 89$, 85 correct and 4 incorrect fragments), and the maximal fragment correct by 90% ($n = 35$, 32 correct and 3 incorrect maximal fragments). We could therefore use a likelihood value threshold above which a fragment or maximal fragment could be tagged as most probably correct. At the other end of the distribution, 90% of maximal fragments with a likelihood value below 8 are wrong ($n = 31$, 3 correct and 28 incorrect maximal fragments). In contrast, small likelihood values do not necessarily indicate that the fragment is false, but rather that the fragment represents a small portion

of the signal. Only below a likelihood value of 0.3, 90% of the fragments are wrong. This implies that using a given likelihood value as cut off to delete fragments would be either inefficient and delete very few fragments, or be inaccurate and delete many correct fragments.

The distribution of ranking values for fragments and maximal fragments are shown in Fig. 7. The left histogram illustrates that the ranking of fragments does not help separating between correct and wrong fragments because the corresponding distributions overlap strongly. However, for the maximal fragments, this overlap is less pronounced (right histogram): correct maximal fragments have a ranking value usually better than 10, while maximal fragments whose ranking is worse than this value are mostly wrong.

These observations can help the user in the identification process by tagging the maximal fragments with a likelihood value above 20 or a ranking value better than 10 as probably correct, and the others as likely wrong. An example of such tagging is shown for CFC-11 in Table 3.

3.4 Performance of the model on the validation set

The validation set (Table 2) is made of 21 compounds, with 457 reconstructed fragments, including 47 maximal fragments. The performance of the reconstructions is similar as with the training set: for 95% of the compounds, more than 90% of the reconstructed signal is correct (Fig. 10). This illustrates that our identification method can be applied to different dataset, while producing similar performances.

3.5 How wrong knapsack solutions are rejected and implications for computation runtime

Figure 8 displays, for each compound of the training and validation sets, how many knapsack fragments are generated, kept and rejected. For compounds where less than ≈ 50 knapsack fragments were generated, we observe that most fragments have gone through the iterating steps of the workflow (see Fig. 1), and are therefore validated (Fig. 8, blue crosses) or deleted, as singletons (grey 'x') or as being below the LOD (mauve stars). On the other hand, for compounds with more than ≈ 100 knapsack fragments, the number of fragments gone through the computation-intensive iterating steps do not exceed ≈ 100 , even if more than 1000 knapsack fragments were generated. In such cases, most fragments are rejected at Step 9 of the workflow (Fig. 1). This behaviour may explain why the computation runtime does not increase linearly with the number of generated knapsack fragments, as discussed hereafter.

Details about the computation runtime are given in Fig. 9. For all compounds, the knapsack step (Step 1) represents only a minor part of the runtime, proving that the optimisation of this computation step is appropriate for the considered compounds. Above 1000 generated knapsack formulae, the graph construction (Step 2) represents an important portion of the runtime, but remains minor for all compounds with less than 1000 knapsack fragments.

For most compounds, the most computation intensive step is the optimisation of multiple isotopocule profiles (Step 7), which uses the machine learning tool `lmfit`. However, the runtime of Step 7 does not increase linearly with the number of knapsack solutions, but for most compounds is limited to less than 10 seconds. We attribute this to the fact that only a limited number of most likely knapsack fragments go through this expensive step.

4 Conclusion

Gas chromatography followed by electron ionisation and high-resolution mass spectrometry is increasingly used for ambient air quality measurements and is a promising technique for non-target screening of trace gases. However, data analysis tools to support automated identification of unknown substances, especially in cases where the molecular ion is absent from the obtained mass spectrum, are still lacking.

In this work, we have developed a novel algorithm to allow reconstruction of the chemical formula based on the measured mass of fragments likely belonging to the same substance. The developed method specifically addresses the observation that peaks with low signal have a higher mass uncertainty, by using a specific uncertainty for each measured mass peak. This approach allows to use all measured data simultaneously, while giving more weight to more accurate mass peaks. In addition, our method does not require the molecular ion to be present, and can still reconstruct the chemical formula of all other detected fragments. This is important for electron-ionisation spectra, where the molecular ion is absent in approx. 40% of the cases.

Overall, the method performs well on very heterogeneous compounds, representing nine different atoms and many different structures, over a large mass domain from 30 m/z to over 300 m/z : for more than 90% of the compounds, more than 90% of the signal has been assigned to the correct chemical formula. The reconstruction becomes less reliable with decreasing number of detected mass peaks, in case of compounds measured at very low molar fraction.

5 Data Availability

A zip version of the Python code is provided for the review process. The Python code will be released upon publication with the acronym ALPINAC, for ALgorithmic Process for Identification of Non-targeted Atmospheric Compounds, under an open-source license at <https://gitlab.empa.ch> and <https://gitlab.inria.fr/guillevi/alpinac/>. In addition, the final published version will be archived on the Zenodo repository.

Additional information

Supplementary information is available in the additional file "[sup_ALPINAC_20210319.pdf](#)".

Competing interests

The authors declare that they have no competing interests.

Author's contributions

MG designed the identification algorithms, with inputs from SR, MKV and PS. AG suggested implementation choices and code improvements. MG and AG wrote the Python software. MKV, PS, MH and MG prepared the standards and collected the ToF data. MG, AG wrote the manuscript with contributions from all co-authors. SR, MKV and LE secured the HALOSEARCH Empa funding.

Acknowledgements

M. Guillevic thanks Luke Western, Stuart Grange and Minsu Kim for valuable discussions related to machine learning. A. Guillevic thanks Paul Zimmermann for discussions on the knapsack problem, and Stephane Glondu for his valuable help in installing Python packages for computational chemistry and setting the development environment on Linux platforms. We thank Georgios Papadopoulos, Mike Cubison (Tofwerk AG) and Harald Stark (Aerodyne Research Inc.) for their support related to the TOF instrument. We thank G. and D. Jones who provided linguistic help. MG is funded by the Empa research grant HALOSEARCH. We acknowledge funding by the DAAMAA project from the Swiss Polar Institute.

Author details

¹Laboratory for Air Pollution /Environmental Technology, Empa, Swiss Federal Laboratories for Materials Science and Technology, Ueberlandstrasse 129, 8600 Dübendorf, Switzerland. ²Université de Lorraine, CNRS, Inria, LORIA, 54000 Nancy, France.

References

1. Acierno, V., Yener, S., Alewijn, M., Biasioli, F., van Ruth, S.: Factors contributing to the variation in the volatile composition of chocolate: Botanical and geographical origins of the cocoa beans, and brand-related formulation and processing. *Food Research International* **84**, 86–95 (2016). doi:[10.1016/j.foodres.2016.03.022](https://doi.org/10.1016/j.foodres.2016.03.022)
2. Hollender, J., Schymanski, E.L., Singer, H.P., Ferguson, P.L.: Nontarget screening with high resolution mass spectrometry in the environment: Ready to go? *Environmental Science & Technology* **51**(20), 11505–11512 (2017). doi:[10.1021/acs.est.7b02184](https://doi.org/10.1021/acs.est.7b02184)
3. Alygizakis, N.A., Samanipour, S., Hollender, J., Ibáñez, M., Kaserzon, S., Kokkali, V., van Leerdam, J.A., Mueller, J.F., Pijnappels, M., Reid, M.J., Schymanski, E.L., Slobodnik, J., Thomaidis, N.S., Thomas, K.V.: Exploring the potential of a global emerging contaminant early warning network through the use of retrospective suspect screening with high-resolution mass spectrometry. *Environmental Science & Technology* **52**(9), 5135–5144 (2018). doi:[10.1021/acs.est.8b00365](https://doi.org/10.1021/acs.est.8b00365)
4. Kruve, A.: Semi-quantitative non-target analysis of water with liquid chromatography/high-resolution mass spectrometry: How far are we? *Rapid Communications in Mass Spectrometry* **33**(S3), 54–63 (2018). doi:[10.1002/rcm.8208](https://doi.org/10.1002/rcm.8208)
5. Sturges, W.T., Wallington, T.J., Hurley, M.D., Shine, K.P., Sihra, K., Engel, A., Oram, D.E., Penkett, S.A., Mulvaney, R., Brenninkmeijer, C.A.M.: A potent greenhouse gas identified in the atmosphere: SF₅CF₃. *Science* **289**(5479), 611–613 (2000). doi:[10.1126/science.289.5479.611](https://doi.org/10.1126/science.289.5479.611)
6. Veenaaas, C., Ripszam, M., Haglund, P.: Analysis of volatile organic compounds in indoor environments using thermal desorption with comprehensive two-dimensional gas chromatography and high-resolution time-of-flight mass spectrometry. *Journal of Separation Science* **43**(8), 1489–1498 (2020). doi:[10.1002/jssc.201901103](https://doi.org/10.1002/jssc.201901103)
7. Vollmer, M.K., Reimann, S., Hill, M., Brunner, D.: First observations of the fourth generation synthetic halocarbons HFC-1234yf, HFC-1234ze(E), and HCFC-1233zd(E) in the atmosphere. *Environmental Science & Technology* **49**(5), 2703–2708 (2015). doi:[10.1021/es505123x](https://doi.org/10.1021/es505123x)
8. Laube, J.C., Newland, M.J., Hogan, C., Brenninkmeijer, C.A.M., Fraser, P.J., Martinerie, P., Oram, D.E., Reeves, C.E., Röckmann, T., Schwander, J., Witrant, E., Sturges, W.T.: Newly detected ozone-depleting substances in the atmosphere. *Nature Geosciences* **7**, 266–269 (2014). doi:[10.1038/ngeo2109](https://doi.org/10.1038/ngeo2109)
9. Ivy, D.J., Arnold, T., Harth, C.M., Steele, L.P., Mühle, J., Rigby, M., Salameh, P.K., Leist, M., Krummel, P.B., Fraser, P.J., Weiss, R.F., Prinn, R.G.: Atmospheric histories and growth trends of C₄F₁₀, C₅F₁₂, C₆F₁₄, C₇F₁₆ and C₈F₁₈. *Atmospheric Chemistry and Physics* **12**(9), 4313–4325 (2012). doi:[10.5194/acp-12-4313-2012](https://doi.org/10.5194/acp-12-4313-2012)
10. Loos, M., Gerber, C., Corona, F., Hollender, J., Singer, H.: Accelerated isotope fine structure calculation using pruned transition trees. *Analytical Chemistry* **87**(11), 5738–5744 (2015). doi:[10.1021/acs.analchem.5b00941](https://doi.org/10.1021/acs.analchem.5b00941)
11. Jaeger, C., Hoffmann, F., Schmitt, C.A., Lise, J.: Automated annotation and evaluation of in-source mass spectra in GC/atmospheric pressure chemical ionization-MS-based metabolomics. *Analytical Chemistry* **88**(19), 9386–9390 (2016). doi:[10.1021/acs.analchem.6b02743](https://doi.org/10.1021/acs.analchem.6b02743)
12. Loos, M., Singer, H.: Nontargeted homologue series extraction from hyphenated high resolution mass spectrometry data. *Journal of Cheminformatics* **9**(12), 1–11 (2017). doi:[10.1186/s13321-017-0197-z](https://doi.org/10.1186/s13321-017-0197-z)
13. Ruttikies, C., Schymanski, E.L., Wolf, S., Hollender, J., Neumann, S.: MetFrag relaunched: incorporating strategies beyond in silico fragmentation. *Journal of Cheminformatics* **8**(3) (2016). doi:[10.1186/s13321-016-0115-9](https://doi.org/10.1186/s13321-016-0115-9)

14. Léon, A., Cariou, R., Hutinet, S., Hurel, J., Guitton, Y., Tixier, C., Munsch, C., Antignac, J.-P., Dervilly-Pinel, G., Le Bizec, B.: HaloSeeker 1.0: A user-friendly software to highlight halogenated chemicals in nontargeted high-resolution mass spectrometry data sets. *Analytical Chemistry* **91**(5), 3500–3507 (2019). doi:[10.1021/acs.analchem.8b05103](https://doi.org/10.1021/acs.analchem.8b05103)
15. Gross, J.H.: *Mass Spectrometry: A Textbook*, 3rd edn. Springer, Berlin Heidelberg (2017). doi:[10.1007/978-3-319-54398-7](https://doi.org/10.1007/978-3-319-54398-7)
16. NIST: NIST EPA NIH Mass Spectral Library. Online database (2020). <https://www.nist.gov/srd/nist-standard-reference-database-1a-v17> Accessed 11.03.2020
17. Schymanski, E.L., Meringer, M., Brack, W.: Automated strategies to identify compounds on the basis of GC/MS and calculated properties. *Analytical Chemistry* **83**(3), 903–912 (2011). doi:[10.1021/ac102574h](https://doi.org/10.1021/ac102574h)
18. Schymanski, E.L., Gallampois, C.M.J., Krauss, M., Meringer, M., Neumann, S., Schulze, T., Wolf, S., Brack, W.: Consensus structure elucidation combining GC/MS, structure generation, and calculated properties. *Analytical Chemistry* **84**(7), 3287–3295 (2012). doi:[10.1021/ac203471y](https://doi.org/10.1021/ac203471y)
19. Obersteiner, F., Bönisch, H., Engel, A.: An automated gas chromatography time-of-flight mass spectrometry instrument for the quantitative analysis of halocarbons in air. *Atmospheric Measurement Techniques* **9**(1), 179–194 (2016). doi:[10.5194/amt-9-179-2016](https://doi.org/10.5194/amt-9-179-2016)
20. Röhler, L., Schlabach, M., Haglund, P., Breivik, K., Kallenborn, R., Bohlin-Nizzetto, P.: Non-target and suspect characterisation of organic contaminants in arctic air – part 2: Application of a new tool for identification and prioritisation of chemicals of emerging arctic concern in air. *Atmospheric Chemistry and Physics* **20**(14), 9031–9049 (2020). doi:[10.5194/acp-20-9031-2020](https://doi.org/10.5194/acp-20-9031-2020)
21. Schuck, T.J., Lefrançois, F., Gallmann, F., Wang, D., Jesswein, M., Hoker, J., Bönisch, H., Engel, A.: Establishing long-term measurements of halocarbons at Taunus Observatory. *Atmospheric Chemistry and Physics* **18**(22), 16553–16569 (2018). doi:[10.5194/acp-18-16553-2018](https://doi.org/10.5194/acp-18-16553-2018)
22. Prinn, R.G., Weiss, R.F., Arduini, J., Arnold, T., DeWitt, H.L., Fraser, P.J., Ganesan, A.L., Gasore, J., Harth, C.M., Hermansen, O., Kim, J., Krummel, P.B., Li, S., Loh, Z.M., Lunder, C.R., Maione, M., Manning, A.J., Miller, B.R., Mitrevski, B., Mühle, J., O'Doherty, S., Park, S., Reimann, S., Rigby, M., Saito, T., Salameh, P.K., Schmidt, R., Simmonds, P.G., Steele, L.P., Vollmer, M.K., Wang, R.H., Yao, B., Yokouchi, Y., Young, D., Zhou, L.: History of chemically and radiatively important atmospheric gases from the advanced global atmospheric gases experiment (AGAGE). *Earth System Science Data* **10**(2), 985–1018 (2018). doi:[10.5194/essd-10-985-2018](https://doi.org/10.5194/essd-10-985-2018)
23. Vollmer, M.K., Young, D., Trudinger, C.M., Mühle, J., Henne, S., Rigby, M., Park, S., Li, S., Guillevic, M., Mitrevski, B., Harth, C.M., Miller, B.R., Reimann, S., Yao, B., Steele, L.P., Wyss, S.A., Lunder, C.R., Arduini, J., McCulloch, A., Wu, S., Rhee, T.S., Wang, R.H.J., Salameh, P.K., Hermansen, O., Hill, M., Langenfelds, R.L., Ivy, D., O'Doherty, S., Krummel, P.B., Maione, M., Etheridge, D.M., Zhou, L., Fraser, P.J., Prinn, R.G., Weiss, R.F., Simmonds, P.G.: Atmospheric histories and emissions of chlorofluorocarbons CFC-13 (CClF₃), Σ CFC-114 (C₂Cl₂F₄), and CFC-115 (C₂ClF₅). *Atmospheric Chemistry and Physics* **18**(2), 979–1002 (2018). doi:[10.5194/acp-18-979-2018](https://doi.org/10.5194/acp-18-979-2018)
24. Prinn, R.G., Weiss, R.F., Fraser, P.J., Simmonds, P.G., Cunnold, D.M., Alyea, F.N., O'Doherty, S., Salameh, P., Miller, B.R., Huang, J., Wang, R.H.J., Hartley, D.E., Harth, C., Steele, L.P., Sturrock, G., Midgley, P.M., McCulloch, A.: A history of chemically and radiatively important gases in air deduced from *ale/gage/agage*. *Journal of Geophysical Research: Atmospheres* **105**(D14), 17751–17792 (2000). doi:[10.1029/2000JD900141](https://doi.org/10.1029/2000JD900141)
25. Arnold, T., Mühle, J., Salameh, P.K., Harth, C.M., Ivy, D.J., Weiss, R.F.: Automated measurement of nitrogen trifluoride in ambient air. *Analytical Chemistry* **84**(11), 4798–4804 (2012). doi:[10.1021/ac300373e](https://doi.org/10.1021/ac300373e). PMID: 22607353
26. Guillevic, M., Vollmer, M.K., Wyss, S.A., Leuenberger, D., Ackermann, A., Pascale, C., Niederhauser, B., Reimann, S.: Dynamic–gravimetric preparation of metrologically traceable primary calibration standards for halogenated greenhouse gases. *Atmospheric Measurement Techniques* **11**(6), 3351–3372 (2018). doi:[10.5194/amt-11-3351-2018](https://doi.org/10.5194/amt-11-3351-2018)
27. Schymanski, E.L., Jeon, J., Gulde, R., Fenner, K., Ruff, M., Singer, H.P., Hollender, J.: Identifying small molecules via high resolution mass spectrometry: Communicating confidence. *Environmental Science & Technology* **48**(4), 2097–2098 (2014). doi:[10.1021/es5002105](https://doi.org/10.1021/es5002105)
28. Miller, B.R., Weiss, R.F., Salameh, P.K., Tanhua, T., Grealley, B.R., Mühle, J., Simmonds, P.G.: Medusa: A sample preconcentration and GC/MS detector system for in situ measurements of atmospheric trace halocarbons, hydrocarbons, and sulfur compounds. *Analytical Chemistry* **80**(5), 1536–1545 (2008). doi:[10.1021/ac702084k](https://doi.org/10.1021/ac702084k). PMID: 18232668
29. United Nations: Montreal protocol on substances that deplete the ozone layer. Montreal, 16 September 1987. Amendment to the Montreal protocol on substances that deplete the ozone layer, Kigali, 15 October 2016. Technical Report C.N.872.2016.TREATIES-XXVII.2.f (Depositary Notification), United Nations, New York, 10017 (2016)
30. Brown, J.S.: Introduction to hydrofluoro-olefin alternatives for high global warming potential hydrofluorocarbon refrigerants. *HVAC&R Research* **19**(6), 693–704 (2013). doi:[10.1080/10789669.2013.802149](https://doi.org/10.1080/10789669.2013.802149)
31. The HDF Group: HDF5 documentation. Online (2020). <https://portal.hdfgroup.org/display/HDF5/HDF5> Accessed 07.12.2020
32. Pap, T.L., Pápai, Z.: Application of a new mathematical function for describing chromatographic peaks. *Journal of Chromatography A* **930**(1), 53–60 (2001). doi:[10.1016/S0021-9673\(01\)01163-3](https://doi.org/10.1016/S0021-9673(01)01163-3)
33. Schroepel, R., Shamir, A.: A $T = O(2^{n/2})$, $S = O(2^{n/4})$ Algorithm for Certain NP-Complete Problems. *SIAM Journal on Computing* **10**(3), 456–464 (1981). doi:[10.1137/0210033](https://doi.org/10.1137/0210033)
34. Odlyzko, A.M.: The rise and fall of knapsack cryptosystems. In: Pomerance, C. (ed.) *Cryptology and Computational Number Theory. Proceedings of Symposia in Applied Mathematics*, August 6–7, 1989, Boulder, Colorado, vol. 42, pp. 75–88. AMS, Providence, Rhode Island (1991). http://www.dtc.umn.edu/~odlyzko/doc/arch/knapsack_survey.pdf. <https://bookstore.ams.org/psapm-42>
35. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*, 3rd edn. The MIT Press,

- Cambridge, Massachusetts, London, England (2009)
36. Meija, J., Coplen, T.B., Berglund, M., Brand, W.A., Bièvre, P.D., Gröning, M., Holden, N.E., Irrgeher, J., Loss, R.D., Walczyk, T., Prohaska, T.: Atomic weights of the elements 2013 (IUPAC technical report). *Pure and Applied Chemistry* **88**(3), 265–291 (2016). doi:[10.1515/pac-2015-0305](https://doi.org/10.1515/pac-2015-0305)
 37. Kind, T., Fiehn, O.: Seven golden rules for heuristic filtering of molecular formulas obtained by accurate mass spectrometry. *BMC Bioinformatics* **8**(105), 1–20 (2007). doi:[10.1186/1471-2105-8-105](https://doi.org/10.1186/1471-2105-8-105)
 38. Dührkop, K., Ludwig, M., Meusel, M., Böcker, S.: Faster mass decomposition. In: Darling, A., Stoye, J. (eds.) *Algorithms in Bioinformatics*, pp. 45–58. Springer, Berlin, Heidelberg (2013). doi:[10.1007/978-3-642-40453-5_5.1307.7805](https://doi.org/10.1007/978-3-642-40453-5_5.1307.7805)
 39. Hagberg, A.A., Schult, D.A., Swart, P.J.: Exploring network structure, dynamics, and function using NetworkX. In: Varoquaux, G., Vaught, T., Millman, J. (eds.) *Proceedings of the 7th Python in Science Conference*, Pasadena, CA USA, pp. 11–15 (2008)
 40. Hufsky, F., Rempt, M., Rasche, F., Pohnert, G., Böcker, S.: De novo analysis of electron impact mass spectra using fragmentation trees. *Analytica Chimica Acta* **739**, 67–76 (2012). doi:[10.1016/j.aca.2012.06.021](https://doi.org/10.1016/j.aca.2012.06.021)
 41. Yergey, J.A.: A general approach to calculating isotopic distributions for mass spectrometry. *International Journal of Mass Spectrometry and Ion Physics* **52**, 337–349 (1983). doi:[10.1002/jms.4498](https://doi.org/10.1002/jms.4498). e4498 JMS-20-0003
 42. Roussis, S.G., Proulx, R.: Reduction of chemical formulas from the isotopic peak distributions of high-resolution mass spectra. *Analytical Chemistry* **75**(6), 1470–1482 (2003). doi:[10.1021/ac020516w](https://doi.org/10.1021/ac020516w)
 43. Newville, M., Stensitzki, T., Allen, D.B., Ingargiola, A.: LMFIT: Non-linear least-square minimization and curve-fitting for Python. Zenodo (2014). doi:[10.5281/zenodo.11813](https://doi.org/10.5281/zenodo.11813). <https://doi.org/10.5281/zenodo.11813>
 44. Kwiecien, N.W., Bailey, D.J., Rush, M.J.P., Cole, J.S., Ulbrich, A., Hebert, A.S., Westphall, M.S., Coon, J.J.: High-resolution filtering for improved small molecule identification via GC/MS. *Analytical Chemistry* **87**(16), 8328–8335 (2015). doi:[10.1021/acs.analchem.5b01503](https://doi.org/10.1021/acs.analchem.5b01503)
 45. Ellison, S.L.R., Williams, A.: *Eurachem/CITAC Guide: Quantifying Uncertainty in Analytical Measurement*, 3rd edn. Eurachem, (2012). Eurachem. ISBN 978-0-948926-30-3. Available from www.eurachem.org
 46. Vollmer, M.K., Rhee, T.S., Rigby, M., Hofstetter, D., Hill, M., Schoenenberger, F., Reimann, S.: Modern inhalation anesthetics: Potent greenhouse gases in the global atmosphere. *Geophysical Research Letters* **42**(5), 1606–1611 (2015). doi:[10.1002/2014GL062785](https://doi.org/10.1002/2014GL062785)

Figures

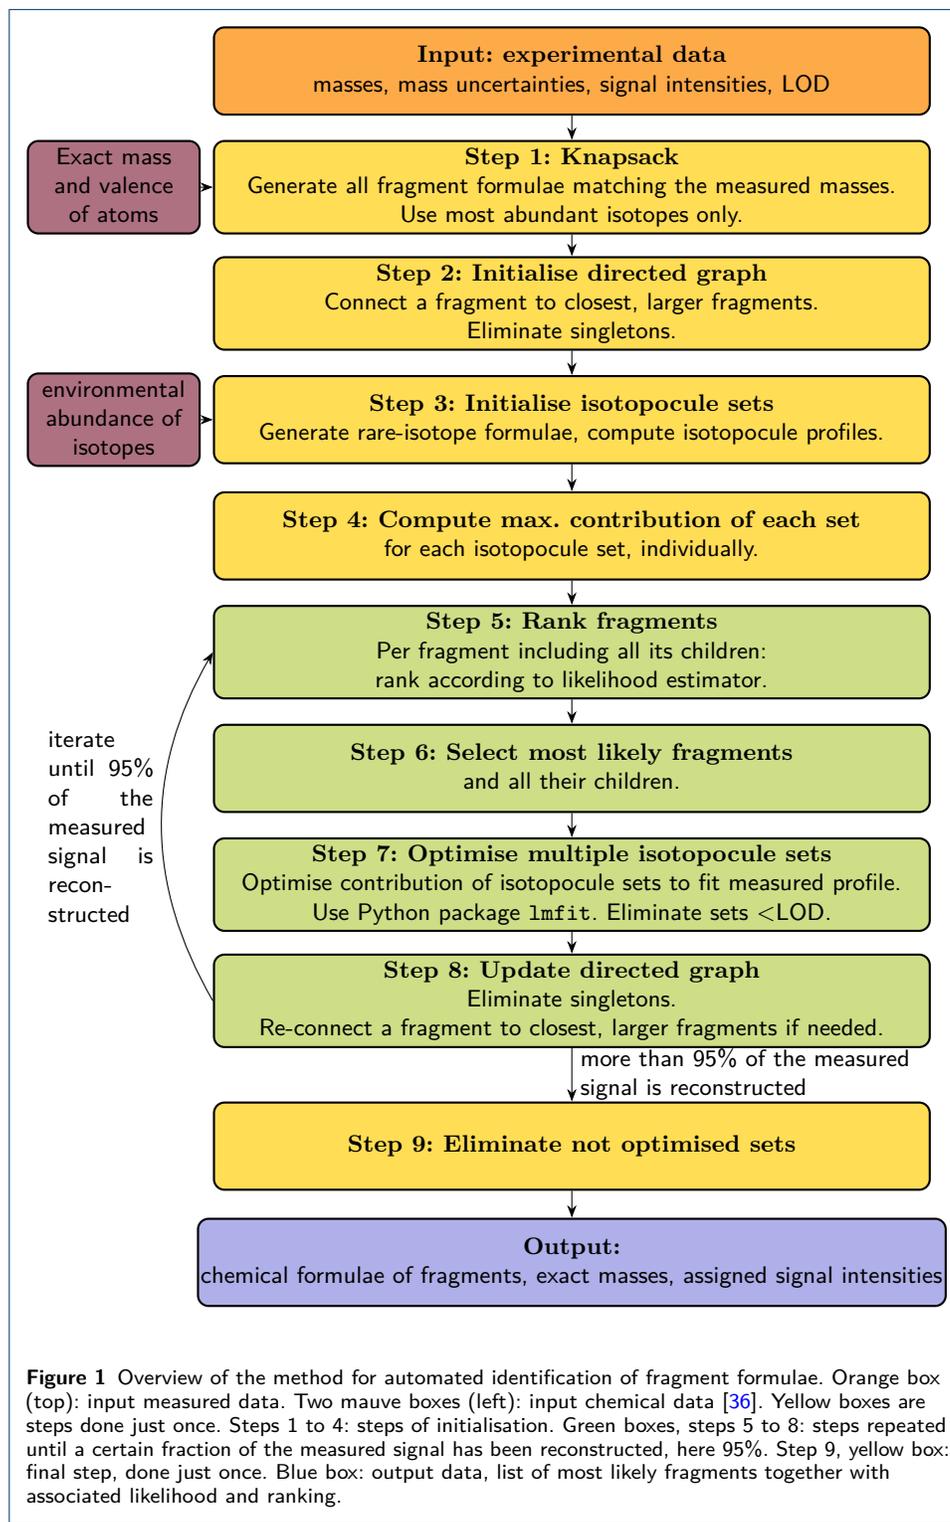
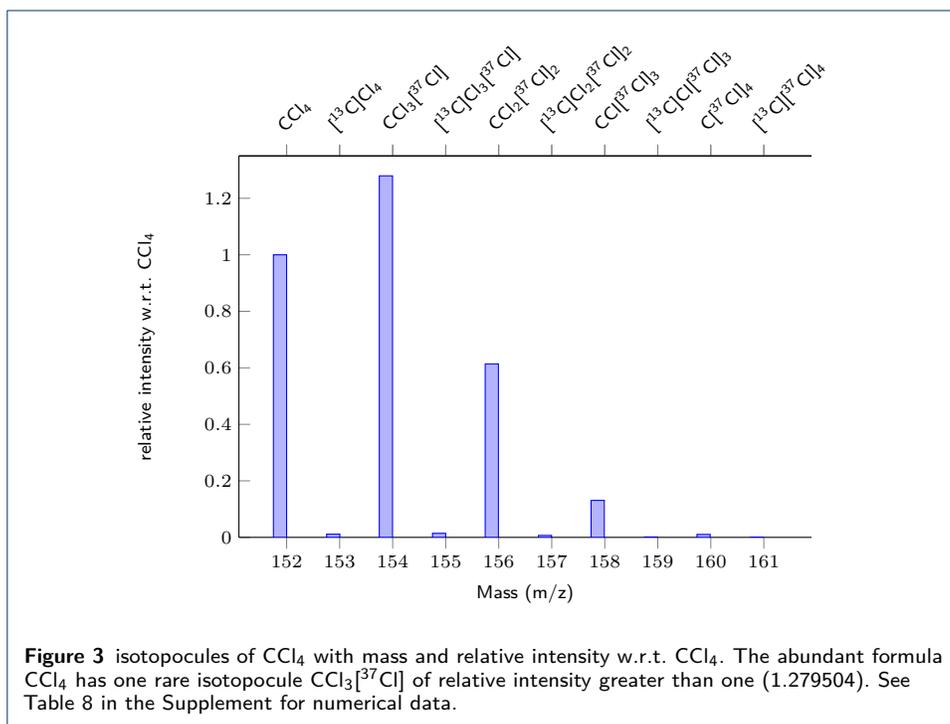
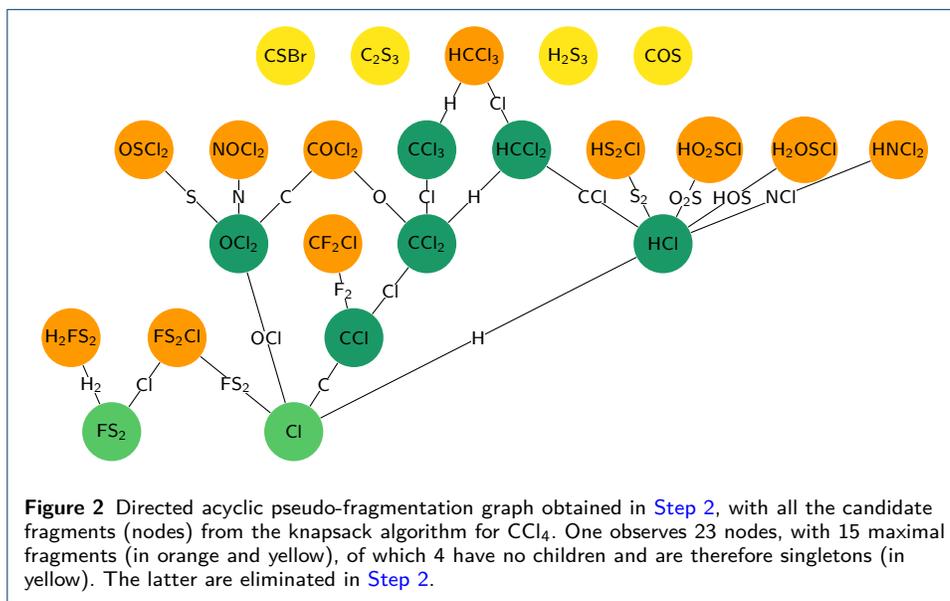
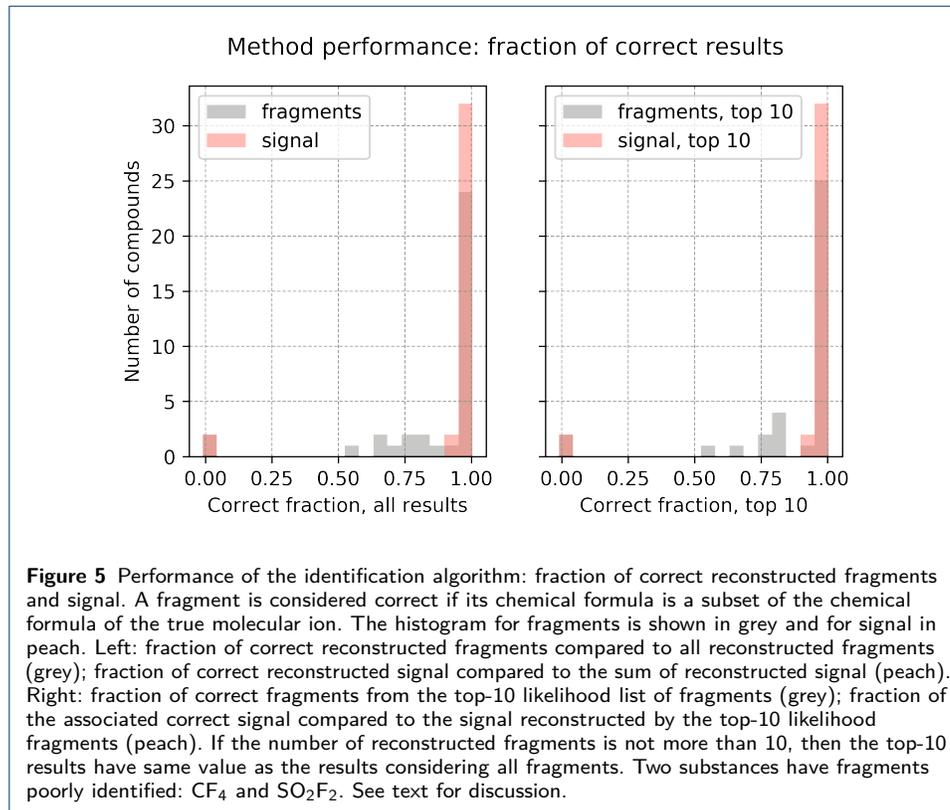
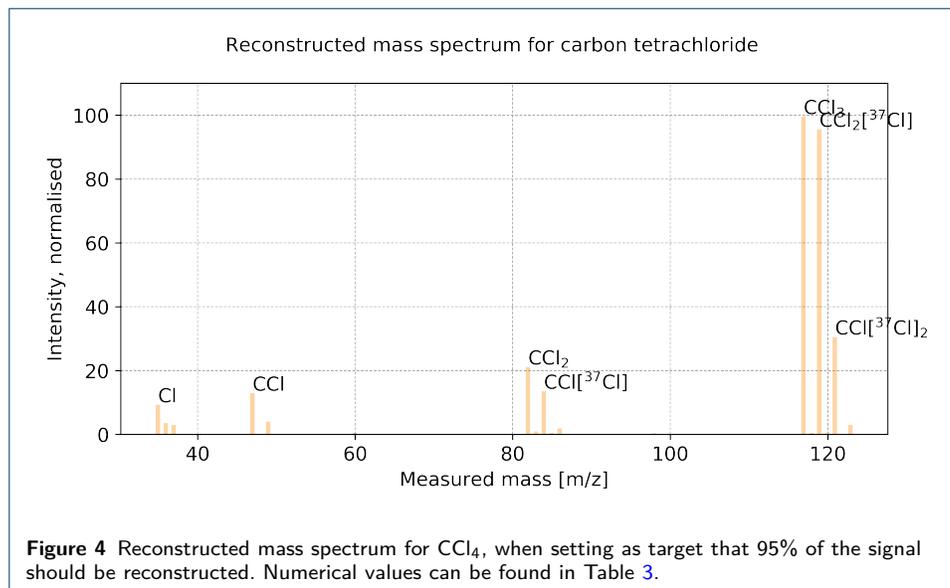
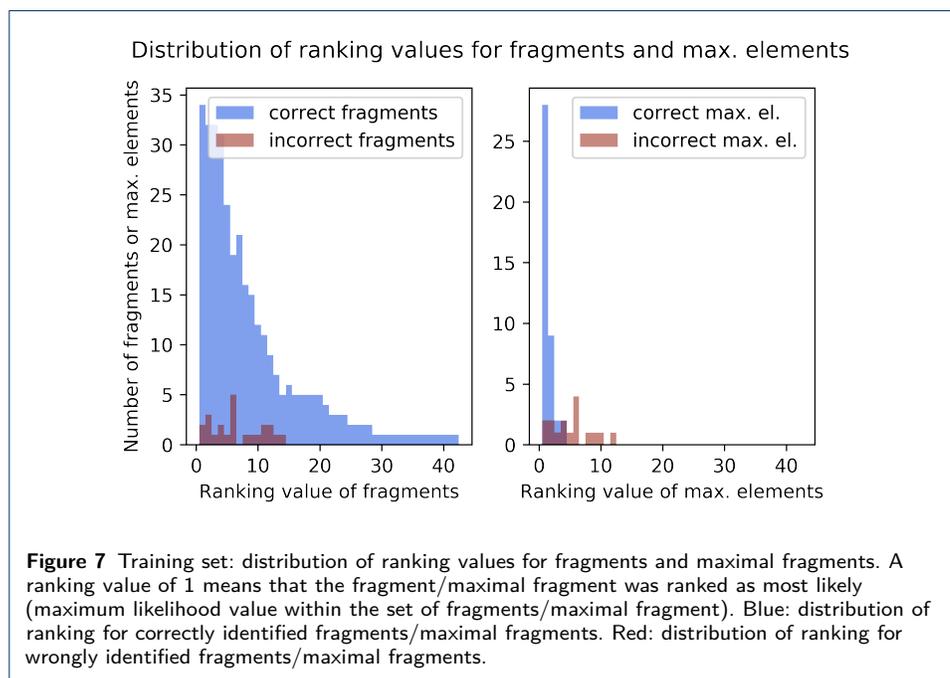
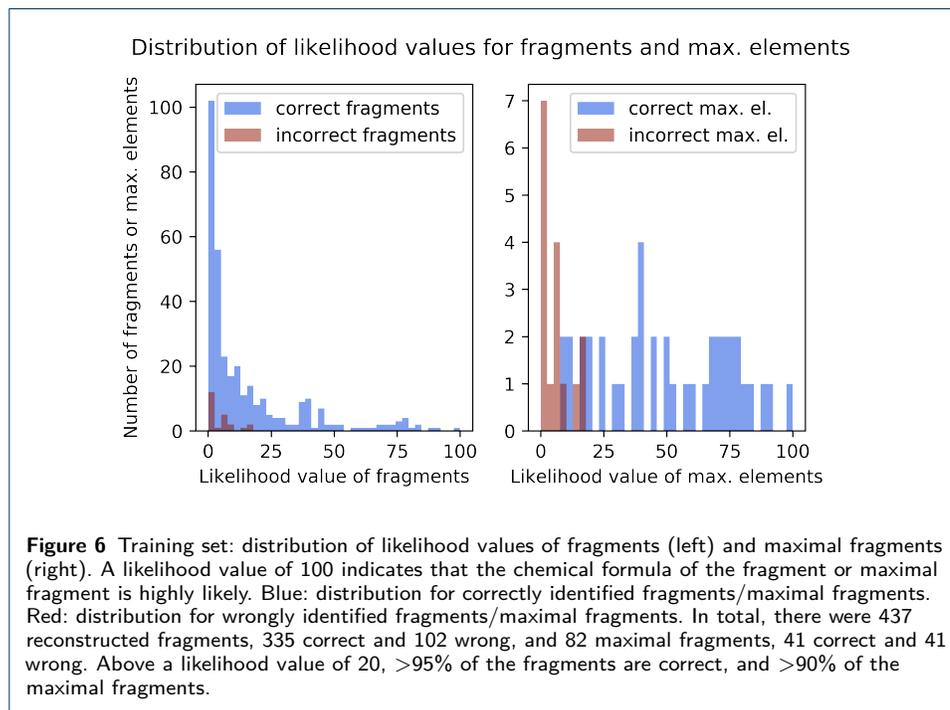
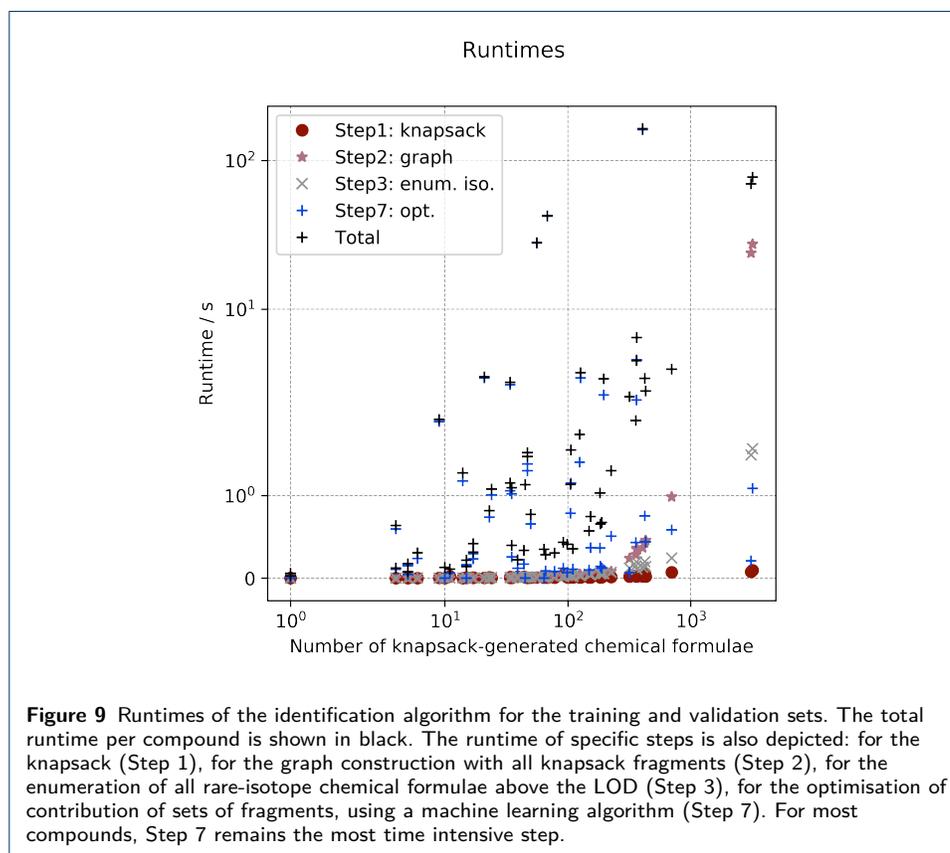
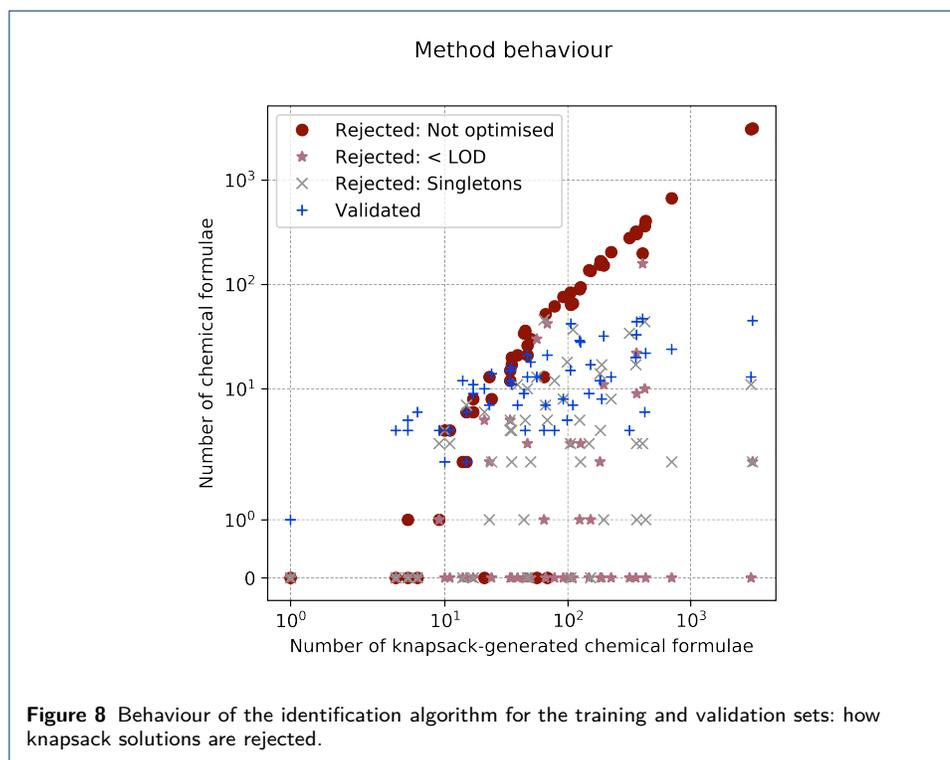


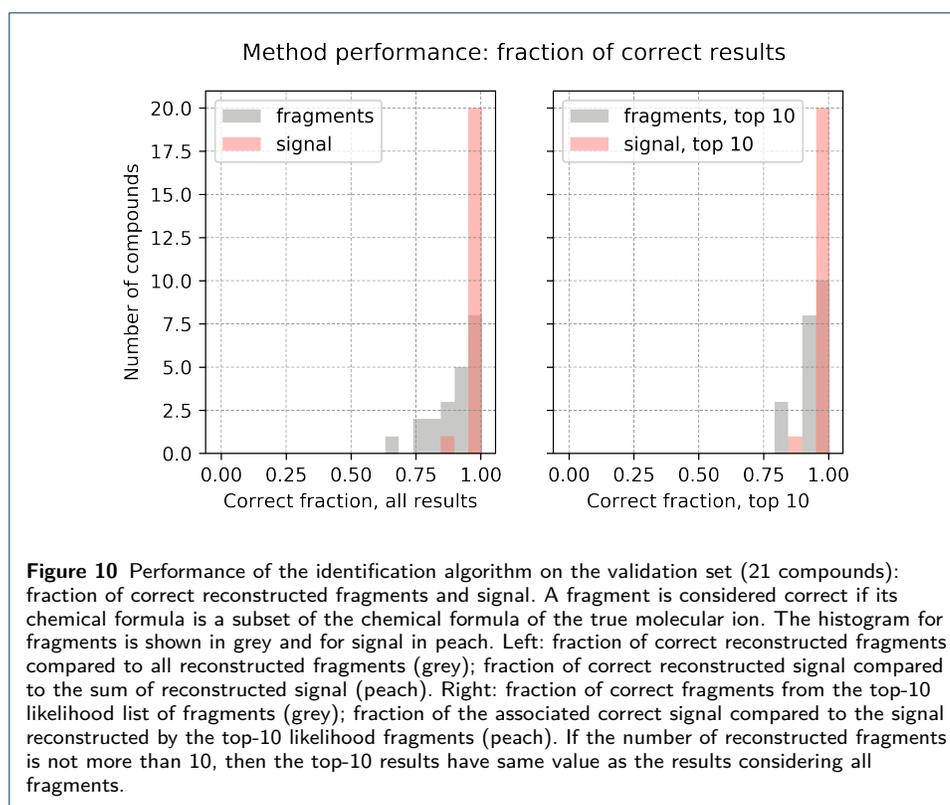
Figure 1 Overview of the method for automated identification of fragment formulae. Orange box (top): input measured data. Two mauve boxes (left): input chemical data [36]. Yellow boxes are steps done just once. Steps 1 to 4: steps of initialisation. Green boxes, steps 5 to 8: steps repeated until a certain fraction of the measured signal has been reconstructed, here 95%. Step 9, yellow box: final step, done just once. Blue box: output data, list of most likely fragments together with associated likelihood and ranking.











Tables

Table 1 Known compounds used as training set. These 36 substances are routinely measured within the AGAGE network [22]. Identification and quantification of these compounds have been done by [23–26]. Present chemical elements are: H, C, N, O, F, S, Cl, Br and I. These are the chemical elements used as input for the knapsack algorithm.

Compound	Chemical name	chemical formula	Monoisotopic molecular mass g mol ⁻¹
C ₂ H ₆	ethane	C ₂ H ₆	30.04695
C ₃ H ₈	propane	C ₃ H ₈	44.06260
CH ₃ Cl	chloromethane	CH ₃ Cl	49.99233
COS	carbonyl sulphide	COS	59.96699
NF ₃	nitrogen trifluoride	NF ₃	70.99828
C ₆ H ₆	benzene	C ₆ H ₆	78.04695
CH ₂ Cl ₂	dichloromethane	CH ₂ Cl ₂	83.95336
HCFC-22	chlorodifluoromethane	HCF ₂ Cl	85.97348
CF ₄	tetrafluoromethane	CF ₄	87.99361
C ₇ H ₈	toluene	C ₇ H ₈	92.06260
CH ₃ Br	bromomethane	CH ₃ Br	93.94181
HCFC-142b	1-chloro-1,1-difluoroethane	H ₃ C ₂ F ₂ Cl	99.98913
SO ₂ F ₂	sulfonyl difluoride	SO ₂ F ₂	101.95871
CFC-13	chlorotrifluoromethane	CF ₃ Cl	103.96406
HCFC-141b	1,1-dichloro-1-fluoroethane	H ₃ C ₂ FCl ₂	115.95958
CHCl ₃	chloroform	CHCl ₃	117.91438
CFC-12	dichlorodifluoromethane	CF ₂ Cl ₂	119.93451
C ₂ HCl ₃	1,1,2-trichloroethene	C ₂ HCl ₃	129.91438
CFC-11	trichlorofluoromethane	CFCl ₃	135.90496
HCFC-124	2-chloro-1,1,1,2-tetrafluoroethane	HC ₂ F ₄ Cl	135.97029
PFC-116	perfluoroethane	C ₂ F ₆	137.99042
CH ₃ I	iodomethane	CH ₃ I	141.92795
SF ₆	sulfur hexafluoride	SF ₆	145.96249
Halon-1301	bromo(trifluoro)methane	CF ₃ Br	147.91355
CCl ₄	tetrachloromethane	CCl ₄	151.87541
CFC-115	1-chloro-1,1,2,2,2-pentafluoroethane	C ₂ F ₅ Cl	153.96087
C ₂ Cl ₄	1,1,2,2-tetrachloroethene	C ₂ Cl ₄	163.87541
Halon-1211	bromochlorodifluoromethane	CF ₂ ClBr	163.88400
CFC-114	1,2-dichloro-1,1,2,2-tetrafluoroethane	C ₂ F ₄ Cl ₂	169.93132
CH ₂ Br ₂	dibromomethane	CH ₂ Br ₂	171.85233
CFC-113	1,1,2-trichloro-1,2,2-trifluoroethane	C ₂ F ₃ Cl ₃	185.90177
PFC-218	perfluoropropane	C ₃ F ₈	187.98723
SF ₅ CF ₃	pentafluoro(trifluoromethyl)sulfur	SF ₅ CF ₃	195.95930
PFC-c318	octafluorocyclobutane	C ₄ F ₈	199.98723
Halon-2402	1,2-dibromo-1,1,2,2-tetrafluoroethane	C ₂ F ₄ Br ₂	257.83029
C ₆ F ₁₄	perfluorohexane	C ₆ F ₁₄	337.97764

Table 2 Known compounds used as validation set. Identification and quantification of these compounds has been done by [7, 26, 46]. Present chemical elements: H, C, F, Cl. Chemical elements used as input for the knapsack algorithm are the same as for the training set: H, C, N, O, F, S, Cl, Br and I.

Compound	Chemical name	Chemical formula	Monoisotopic molecular mass g mol ⁻¹
Kigali Amendment to the Montreal Protocol			
HFC-41	fluoromethane	CH ₃ F	34.021878
HFC-32	difluoromethane	CH ₂ F ₂	52.012456
HFC-152	1,2-difluoroethane	C ₂ H ₄ F ₂	66.028106
HFC-152a	1,1-difluoroethane	C ₂ H ₄ F ₂	66.028106
HFC-23	fluoroform	CHF ₃	70.003035
HFC-143	1,1,2-trifluoroethane	C ₂ H ₃ F ₃	84.018685
HFC-143a	1,1,1-trifluoroethane	C ₂ H ₃ F ₃	84.018685
HFC-134	1,1,2,2-tetrafluoroethane	C ₂ H ₂ F ₄	102.009263
HFC-134a	1,1,1,2-tetrafluoroethane	C ₂ H ₂ F ₄	102.009263
HFC-125	pentafluoroethane	C ₂ HF ₅	119.999841
HFC-245ca	1,1,2,2,3-pentafluoropropane	C ₃ H ₃ F ₅	134.015491
HFC-245fa	1,1,1,3,3-pentafluoropropane	C ₃ H ₃ F ₅	134.015491
HFC-365mfc	1,1,1,3,3-pentafluorobutane	C ₄ H ₅ F ₅	148.031141
HFC-236cb	1,1,1,2,2,3-hexafluoropropane	C ₃ H ₂ F ₆	152.006069
HFC-236ea	1,1,1,2,3,3-hexafluoropropane	C ₃ H ₂ F ₆	152.006069
HFC-236fa	1,1,1,3,3,3-hexafluoropropane	C ₃ H ₂ F ₆	152.006069
HFC-227ea	1,1,1,2,3,3,3-heptafluoropropane	C ₃ HF ₇	169.996647
HFC-43-10mee	1,1,1,2,2,3,4,5,5,5-decafluoropentane	C ₅ H ₂ F ₁₀	251.999682
HFOs			
HFO-1234yf	2,3,3,3-tetrafluoroprop-1-ene	H ₂ C ₃ F ₄	114.009263
HFO-1234ze(E)	(E)-1,3,3,3-tetrafluoroprop-1-ene	H ₂ C ₃ F ₄	114.009263
HCFO-1233zd(E)	(E)-1-chloro-3,3,3-trifluoro prop-1-ene	H ₂ C ₃ F ₃ Cl	129.979712

Table 3 Results of the identification algorithm for CFC-11: final list of fragments. 1st column: chemical formula of the fragment, containing only abundant atoms. * indicates that the fragment is not part of the molecular ion. 2nd col.: calculated exact mass of ionised fragment. 3rd col.: percent signal of the fragment and its isotopocules relative to the total measured signal. 4th col.: percent signal of the fragment, all its sub-fragments and all associated isotopocules relative to the total measured signal, computed from Eq. (5). 5th col.: likelihood value computed from Eq. (6). 6th col.: ranking according to decreasing likelihood value. In this example, CFC₁₁ ranked first is the molecular ion.

Fragment	Exact mass	% Assigned signal of the fragment	% Assigned signal of all (sub)fragments	Likelihood	Ranking	Maximal fragment?	Tag
CFC₁₁⁺	135.90496	0.0	98.1	88.3	1	yes	Most likely correct max. fragment
CF ₂ Cl ⁺	100.93611	81.1	95.9	83.9	2	no	
CFCl ⁺	65.96726	5.7	12.1	12.1	3	no	
*HCFCl ⁺	66.97508	0.0	12.3	9.5	4	yes	
CCl ₃ ⁺	116.90656	2.2	10.1	8.1	5	no	
CCl ₂ ⁺	81.93771	2.6	7.9	5.9	6	no	
CCl ⁺	46.96885	3.0	5.3	5.3	7	no	
*HCl ⁺	35.97668	0.1	2.4	2.4	8	no	
FCl ⁺	53.96726	0.0	2.3	2.3	9	no	
Cl ⁺	34.96885	2.3	2.3	2.3	10	no	
CF ⁺	30.99840	1.1	1.1	1.1	11	no	