

Competitive fragmentation modeling of ESI-MS/MS spectra for putative metabolite identification

Felicity Allen · Russ Greiner · David Wishart

Received: 10 March 2014 / Accepted: 14 May 2014 / Published online: 5 June 2014
© Springer Science+Business Media New York 2014

Abstract Electrospray tandem mass spectrometry (ESI-MS/MS) is commonly used in high throughput metabolomics. One of the key obstacles to the effective use of this technology is the difficulty in interpreting measured spectra to accurately and efficiently identify metabolites. Traditional methods for automated metabolite identification compare the target MS or MS/MS spectrum to the spectra in a reference database, ranking candidates based on the closeness of the match. However the limited coverage of available databases has led to an interest in computational methods for predicting reference MS/MS spectra from chemical structures. This work proposes a probabilistic generative model for the MS/MS fragmentation process, which we call competitive fragmentation modeling (CFM), and a machine learning approach for learning parameters for this model from MS/MS data. We show that CFM can be used in both a MS/MS spectrum prediction task (ie, predicting the mass spectrum from a chemical structure), and in a putative metabolite identification task (ranking possible structures for a target MS/MS spectrum). In the MS/MS spectrum prediction task, CFM shows significantly improved performance when compared to a full enumeration of all peaks corresponding to substructures of the molecule. In the metabolite identification task, CFM obtains substantially better rankings for the correct candidate than existing methods (MetFrag and FingerID) on

tripeptide and metabolite data, when querying PubChem or KEGG for candidate structures of similar mass.

Keywords Tandem mass spectrometry · MS/MS · Metabolite identification · Machine learning

1 Introduction

Liquid chromatography combined with electrospray ionisation mass spectrometry (ESI-MS) is one of the most frequently used approaches for conducting metabolomics experiments (Dunn and Ellis 2005; Tautenhahn et al. 2012; Kind and Fiehn 2010; Wishart 2011). Collision-induced dissociation (CID) is usually employed within this procedure, intentionally fragmenting molecules into smaller parts to examine their structure. This is called MS/MS or tandem mass spectrometry. A significant bottleneck in such experiments is the interpretation of the resulting spectra to identify metabolites.

Widely used methods for putative metabolite identification (Sumner et al. 2007), using mass spectrometry, compare a collected MS or MS/MS spectrum for an unknown compound against a database containing reference MS or MS/MS spectra (Stein and Scott 1994; Scheubert et al. 2013; Tautenhahn et al. 2012). Unfortunately, current reference databases are still fairly limited, especially in the case of ESI-MS/MS. At the time of writing, the public Human Metabolome Database (Wishart et al. 2013) contains ESI-MS/MS data for around 800 compounds, which represents only a small fraction of the 40,468 known human metabolites it lists. The publicly available Metlin database (Smith et al. 2005) provides ESI-MS/MS spectra for 11,209 of the 75,000 endogenous and exogenous metabolites it contains, although more than half

Electronic supplementary material The online version of this article (doi:10.1007/s11306-014-0676-4) contains supplementary material, which is available to authorized users.

F. Allen (✉) · R. Greiner · D. Wishart
Department of Computing Science, University of Alberta,
2-21 Athabasca Hall, Edmonton, AB T6G 2E8, Canada
e-mail: felicity.allen@ualberta.ca

of those spectra are for enumerated tripeptides. The public repository MassBank (Horai et al. 2010) contains a more diverse dataset of 31,000 spectra collected on a variety of different instruments, including ESI-MS/MS spectra for ~2,000 unique compounds. However, set against the more than 19 million chemical structures in the Pubchem Compound database (Bolton et al. 2008), an estimated 200,000 plant metabolites (Fiehn 2002), or even the 32,801 manually annotated entries in the database of Chemical Entities of Biological Interest (ChEBI) (Hastings et al. 2013), we see that MS/MS coverage still falls far short of the vast number of known metabolites and molecules of interest.

Consequently, there is substantial interest in finding alternative means for identifying metabolites for which no reference spectra are available (Scheubert et al. 2013). For these cases, one approach to metabolite identification involves first predicting the MS or MS/MS spectrum for each candidate compound from its chemical structure (Heinonen et al. 2008; Wolf et al. 2010; Lindsay et al. 1980; Gasteiger et al. 1992). The interpreter then uses these predicted spectra in place of reference spectra, and labels the target spectrum as the metabolite whose predicted spectrum is the closest match, according to some similarity criteria. A wide range of similarity criteria have been proposed, from weighted counts of the number of matching peaks (Stein and Scott 1994), to more complex probability based measures (Mylonas et al. 2009; Oberacher et al. 2009).

The upshot of this predictive approach is that only a list of candidate molecules is needed, rather than a complete database of reference spectra. However, the restriction to a list of candidate molecules means that this approach still falls short of *de novo* identification of 'unknown unknowns' (Wishart et al. 2009), i.e. we cannot identify molecules not in the list.

The concept of computer-based MS prediction has been around since the Dendral project in the 1960s, when investigators attempted to predict electron ionization (EI) mass spectra using early machine learning methods (Lindsay et al. 1980). More recent approaches to this problem have generally taken one of two forms: rule-based or combinatorial.

Commercial packages, such as Mass Frontier (Thermo Scientific, www.thermoscientific.com), and MS Fragmenter (ACD Labs, www.acdlabs.com), are rule-based, using thousands of manually curated rules to predict fragmentations. Primarily developed for EI fragmentation, these packages have been extended for use with ESI. This current work does not compare against these methods empirically, however in at least one study they have been found to have been out-performed by MetFrag (Wolf et al. 2010), to which we do compare. MOLGEN-MS (Kerber et al. 2006) also applies rule-based fragmentations in

combination with an isotope-dependent matching criteria to rank candidate molecules for a given EI spectrum. Another knowledge-based approach, called MASSIMO, combines chemical knowledge with data; using logistic regression to predict fragmentation probabilities for a particular class of EI fragmentations (Gasteiger et al. 1992).

The other class of algorithms applies a combinatorial fragmentation procedure, enumerating all possible fragments of the original structure by systematically breaking bonds (Hill and Mortishire-Smith 2005; Heinonen et al. 2008; Wolf et al. 2010). First proposed by Hill and Mortishire-Smith (2005), this method has been incorporated into the freely available programs FiD (Heinonen et al. 2008) and MetFrag (Wolf et al. 2010). Both identify the given spectrum with the metabolite that has the most closely matching peaks via such a combinatorial fragmentation. These programs also employ several heuristics in their scoring protocols to emphasise the importance of more probable fragmentations. FiD uses an approximate measure of the dissociation energy of the broken bond, combined with a measure of the energy of the product ion. MetFrag incorporates a similar measure of bond energy combined with a bonus if the neutral loss formed is one of a common subset.

An alternative method, FingerID (Heinonen et al. 2012), takes advantage of the increasing number of available MS/MS spectra, by applying machine learning methods to this task. This program uses support vector machines (SVMs) to predict a chemical fingerprint directly from an MS/MS spectrum, and then searches for the metabolite that most closely matches that predicted fingerprint. For a more extensive review of existing computational methods in MS-based metabolite identification, see Hufsky et al. (2014).

The main problem with the current combinatorial methods is that, while they have very good recall, explaining most if not all peaks in each spectrum, they also have poor precision, predicting many more peaks than are actually observed. MetFrag and FiD attempt to address this problem by adding the heuristics described above. In our work, we investigate an alternative machine learning approach that aims to improve the precision of such combinatorial methods.

We propose a method for learning a generative model of the CID fragmentation process from data. This model estimates the likelihood of any given fragmentation event occurring, thereby predicting those peaks that are most likely to be observed. We hypothesise that increasing the precision of the predicted spectrum in this way will improve our system's ability to accurately identify metabolites. In a similar spirit, Kangas et al. (2012) proposed a machine learning approach for obtaining bond

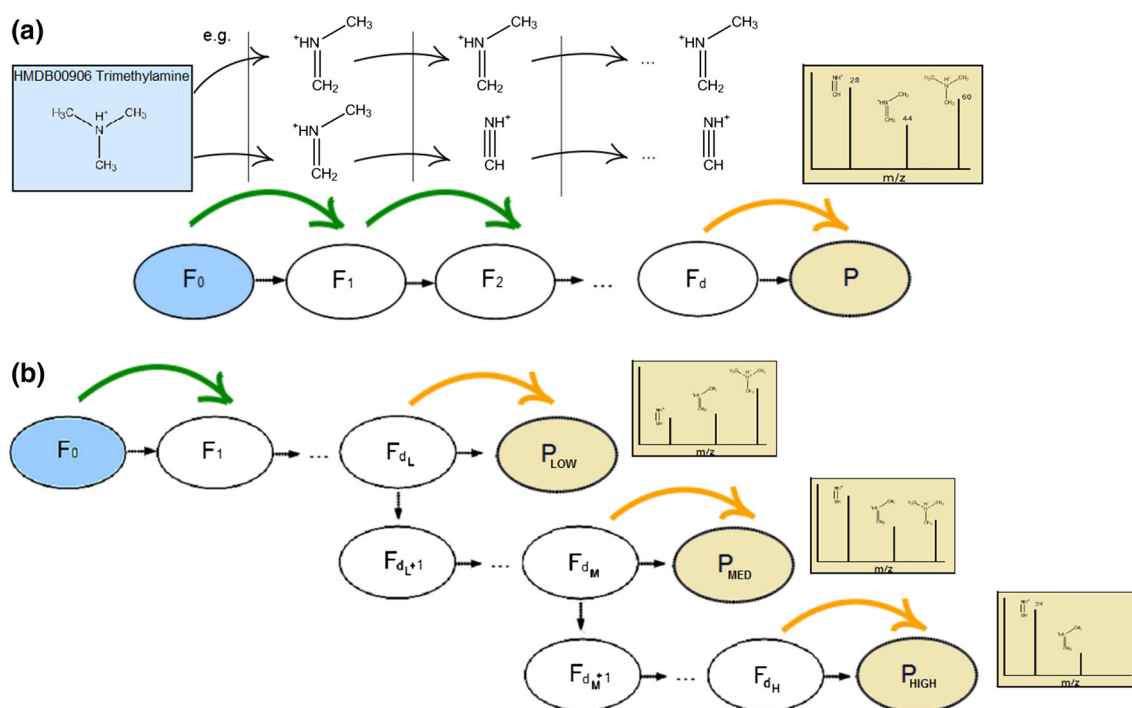


Fig. 1 **a** Single energy competitive fragmentation model (SE-CFM): a stochastic, Markov process of state transitions between charged fragments. **b** Combined energy competitive fragmentation model

(CE-CFM): an extension of SE-CFM that combines information from multiple collision energy spectra into one model

dissociation energies for lipids. Their method uses a different model and training paradigm which, to the authors' knowledge, has not yet been applied to general classes of metabolites.

Section 2 provides details of our proposed model and the training method. Section 3 then reports the experimental results. We will assume the reader knows the foundations of ESI-MS/MS; for an introduction to this process, see de Hoffman and Stroobant (2007).

2 Methods

This section presents our model for the ESI-MS/MS CID fragmentation process, which we call competitive fragmentation modeling (CFM), and a method for deriving parameters for this model from existing MS/MS data. Section 2.1 describes the simplest form of this method; single energy competitive fragmentation modeling (SE-CFM). Section 2.2 then presents an extension of this method, combined energy competitive fragmentation modeling (CE-CFM), which aims to make better use of CID MS/MS spectra measured at different energy levels for the same compound.

Windows executables, cross-platform source code and the trained models used in Sect. 3 are freely available at <http://sourceforge.net/projects/cfm-id/>. A web server interface is

also provided at <http://cfmid.wishartlab.com>. This provides access to the SE-CFM model trained on the Metlin Metabolite data as used in Sect. 3, along with examples of predicted spectra.

2.1 Single energy CFM (SE-CFM)

In single energy CFM (SE-CFM), we model ESI-MS/MS fragmentation as a stochastic, homogeneous, Markov process (Cappé et al. 2005) involving state transitions between charged fragments, as depicted in Fig. 1a.

More formally, the process is described by a fixed length sequence of discrete, random fragment states F_0, F_1, \dots, F_d , where each F_i takes a value from the state space $\mathcal{F} := \{f_1, f_2, \dots, f_{|\mathcal{F}|}\}$, the set of all possible fragments; this state space will be further described in Sect. 2.1.1. A transition model defines the probabilities that each fragment leads to another at one step in the process; see Sect. 2.1.2. An observation model maps the penultimate node F_d to a peak P , which takes on a value in \mathbb{R} that represents the m/z value of the peak to which the final fragment will contribute; see Sect. 2.1.4.

SE-CFM is a latent variable model in which the only observed variables are the initial molecule F_0 and the output peak P ; the fragments themselves are never directly observed. Each output P adds only a small contribution to a single peak in the mass spectrum. In order to predict a

complete mass spectrum, we can run the model forward multiple times to compute the marginal distribution of P .

2.1.1 Fragment state space

We make the following assumptions about the CID fragmentation process. Further details for the motivations of each are provided below, but these generally involve a trade-off between accurately modeling the process and keeping the model computationally tractable.

1. All input molecules have a single positive charge and exist in their most common isotopic form.
2. In a collision, each molecule will break into two fragments.
3. No mass or charge is lost. One of the two fragments must have a single positive charge and the other must be neutral. Combined, the two must contain all the components of the original charged molecule, i.e. all the atoms and electrons.
4. No further sigma bonds can be removed or added during a break, except those connecting hydrogens—i.e. the edges in the molecular graph must remain the same.
5. Rearrangement of pi bonds is allowed and hydrogen atoms may move anywhere in the two resulting fragments, on the condition that both fragments satisfy all valence rules, and standard bond limitations are met—e.g. no bond orders higher than triple.
6. The even electron rule is always satisfied—i.e. no radicals.

Assumption 1 is reasonable as we assume that the first phase of MS/MS successfully restricts the mass range of interest to include only the $[M+H]^+$ precursor ion containing the most abundant isotopes. Since this ion has only a single positive charge, we can safely assume that no multiply-charged ions will be formed in the subsequent MS2 phase. Ensuring that valid $[M+H]^+$ precursor ions are selected in MS1 is beyond the scope of this work; see Katajamaa and Oresic (2007) for a summary of MS1 data processing methods.

Assumptions 2, 4 and 6 do not necessarily hold in real-world spectra (Galezowska et al. 2013; Levsen et al. 2007). However including them substantially reduces the branching factor of the fragment enumeration, making the computations feasible. Since these assumptions do appear to hold in the vast majority of cases, we expect that including them should have minimal negative impact on the experimental results. Note that most 3-way fragmentations can be modeled by two sequential, 2-way fragmentations, so including Assumption 2 should not impact our ability to model most fragmentation events.




Fig. 2 An abstract example of a fragmentation graph, showing a directed acyclic graph of all possible ways in which a particular charged molecule may break to produce smaller charged fragments

Assumption 5 allows for McLafferty Rearrangement and other known fragmentation mechanisms (McLafferty and Turecek 1993).

Our method for enumerating fragments is similar in principle to the combinatorial approach used in MetFrag and FiD (Wolf et al. 2010; Heinonen et al. 2008), with some additional checks to enforce the above assumptions. We systematically break all non-ring bonds in the molecule (excluding those connecting to hydrogens) and all pairs of bonds within each ring. We do this one break at a time, enumerating a subset of fragments with all possible masses that may form after each break, allowing for hydrogen rearrangements. This subset is found by determining the number of additional electrons that can be allocated to either side of the break using integer linear programming to enforce bond constraints—e.g. breaking the middle bond in CCC[CH4+] (SMILES format) gives possible fragments C=[CH3+] (mass = 29.04 Da, loss CC) and C[CH4+] (mass = 31.05 Da, loss C=C), whereas it is not possible to break the triple bond in C#[CH2+] because there is nowhere for the electrons from the bond to go.

The fragmentation procedure is applied recursively on all the produced fragments, to a maximum depth. The result is a directed acyclic graph (DAG) containing all possible charged fragments that may be generated from that molecule. An abstract example of such a fragmentation graph is provided in Fig. 2. Note that for each break, one of the two produced fragments will have no charge. Since it is not possible for a mass spectrometer to detect neutral molecules, we do not explicitly include the neutral fragments in the resulting graph, nor do we recur on their

possible breaks. However neutral loss information may be included on the edges of the graph, indicating how a particular charged fragment was determined. This representation of the fragmentation possibilities as a DAG is similar to that proposed by Böcker and Rasche (2008) with the exception that their nodes contain molecular formulae rather than structures for the ions.

2.1.2 Transition model

Our parametrized transition model assigns a conditional probability to each fragment given the previous fragment in the sequence F_0, F_1, \dots, F_d . Recall that F_t denotes the random fragment state at time t , whereas f_i denotes the i th fragment in the space of all fragments. In the case where f_i has f_j as a possible child fragment in a fragmentation graph, our model assigns a positive probability to the transition from $F_t = f_i$ to $F_{t+1} = f_j$. Furthermore, self-transitions are always allowed, i.e. the probability of transitioning from $F_t = f_i$ to $F_{t+1} = f_i$ is always positive (for the same f_i). We assign 0 probability to all other transitions, i.e. those that are not self-transitions, and that do not exist within any fragmentation graph.

Although the set of possible charged fragments \mathcal{F} is large, the subset of child fragments originating from any particular fragment is relatively small. For example, the requirement that a feasible child fragment must contain a subset of the atoms in the parent fragment rules out many possibilities. Consequently most transitions will be assigned a probability of 0. Note that the assigned probabilities of all transitions originating at a particular fragment, including the self-transition, must sum to one.

We now discuss how we parametrize our transition model. A natural parametrization would be to use a transition matrix containing a separate parameter for every possible fragmentation $f_i \rightarrow f_j$. Unfortunately, we lack sufficient data to learn parameters for every individual fragmentation in this manner. Instead, we look for methods that can generalize by exploiting the tendency of similar molecules to break in similar ways.

2.1.3 Break tendency

We introduce the notion of *break tendency*, which we represent by a value $\theta \in \mathbb{R}$ for each possible fragmentation $f_i \rightarrow f_j$ that models how likely a particular break is to occur. Those fragmentations that are more likely to occur are assigned a higher break tendency value, and those that are less likely are given lower values. We then employ a softmax function to map the break tendencies for all breaks involving a particular parent fragment to probabilities, as defined in Eq. 1 below. This has the effect of capturing the




Fig. 3 Two similar breaks, both resulting in an H_2O neutral loss. The *right case* should be assigned a higher probability, as in the *left case*, the NH_3 is also likely to break away, reducing the probability of the H_2O loss

competition that occurs between different possible breaks within the same molecule. For example, consider the two fragmentations in Fig. 3. Here, although both fragmentations involve an H_2O neutral loss, in the left-hand case, the H_2O loss must compete with the loss of an ammonia group, whereas in the right hand case, it does not. Hence our model might assign an equal break tendency to both cases, but this would still result in a lower probability of fragmentation in the former case, due to the competing ammonia.

We model the probability of a particular break $f_i \rightarrow f_j$ occurring as a function of its break tendency value $\theta_{i,j}$ and that of all other competing breaks from the same parent, as follows:

$$\rho(f_i, f_j) = \begin{cases} \frac{\exp \theta_{i,j}}{1 + \sum_k \exp \theta_{i,k}} & : f_i \neq f_j \text{ and } f_i \rightarrow f_j \text{ is possible} \\ \frac{1}{1 + \sum_k \exp \theta_{i,k}} & : f_i = f_j \\ 0 & : f_i \rightarrow f_j \text{ is not possible} \end{cases} \quad (1)$$

where the sums iterate over all k for which $f_i \rightarrow f_k$ is possible.

Since the break tendency is a relative measure, it makes sense to tie it to some reference point. For the purposes of this model, we have assigned the break tendency for a self-transition (i.e. no break occurring) to $\theta_{i,i} = 0$, which gives $\exp \theta_{i,i} = 1$ as shown in (1).

Incorporating chemical features We need to compute $\theta_{i,j}$ for $i \neq j$. To do this we first define a binary feature vector $\Phi_{i,j}$ to describe the characteristics of a given break $f_i \rightarrow f_j$. Such features might include the presence of a particular atom adjacent to the broken bond, or the formation of a specific neutral loss molecule—e.g. see Sect. 3.2. We then use these features to assign a break tendency value using a linear function parameterized by a vector of weights $w \in \mathbb{R}^n$ —i.e. $\theta_{i,j} := w^T \Phi_{i,j}$. This can then be substituted into (1) to generate the probability of transition $f_i \rightarrow f_j$. The first feature of $\Phi_{i,j}$ is a bias term, set to 1 for all breaks. Note that the vector w constitutes the parameters of the CFM model that we will be learning.

2.1.4 Observation model

We model the conditional probability of P using a narrow Gaussian distribution centred around the mass¹ of F_d , i.e. $P|F_d \sim \mathcal{N}(\text{mass}(F_d), \sigma^2)$. The value for σ can be set according to the mass accuracy of the mass spectrometer used. So, we define this observation function to be the following

$$g(m, F_d; \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\frac{m - \text{mass}(F_d)}{\sigma}\right)^2\right\}. \quad (2)$$

Our investigation (see supplementary data) of the mass error of the precursor ions in the Metlin metabolite data used in Sect. 3 found that the distribution of mass errors had a mean offset of ~ 1 ppm, and a narrower shape than a Gaussian distribution. However, in order to model a more general mass error, not specific to a particular instrument or set of empirical data, we think the Gaussian distribution is a reasonable approach.

2.1.5 Selecting parameter values

Our system estimates the values for the parameters w of the proposed model by applying a training procedure to a set of molecules $\mathcal{X} = \{x_1, x_2, \dots, x_{|\mathcal{X}|}\}$, for which we have both the chemical structure and a measured MS/MS spectrum.

For the purposes of this work, we assume we have a measured low, medium and high energy CID MS/MS spectrum for each molecule, which we denote $S(x) = (s_L(x), s_M(x), s_H(x)) \quad \forall x \in \mathcal{X}$. Each spectrum is further defined to be a set of peaks, where each peak is a pair (m, h) , composed of a mass $m \in \mathbb{R}$ and a height (or intensity) $h \in [0, 100] \subset \mathbb{R}$. Note that each spectrum is normalized, such that the peak heights sum to 100.

For this single energy version of the model, we derive parameters for a completely separate model for each of the three energy levels, using data from that level only. Note that if we had data for only one energy level, we could use this method to train a model using just that energy. However Sect. 2.2 will extend this model to combine the three energy spectra for use in a single model. Until then, we will use $s(x)$ to denote whichever of $s_L(x)$, $s_M(x)$ or $s_H(x)$ we are currently considering.

Maximum likelihood We use a maximum likelihood approach for parameter estimation. The likelihood of the data \mathcal{X} , given the parameters w , and incorporating the previously defined transition function ρ and observation function g , is given by

$$\mathcal{L}(w, \mathcal{X}) = \prod_{x \in \mathcal{X}} \prod_{(m, h) \in s(x)} \left(\sum_{F_1 \in C'(x)} \rho(x, F_1; w) \sum_{F_2 \in C'(F_1)} \rho(F_1, F_2; w) \dots \sum_{F_d \in C'(F_{d-1})} \rho(F_{d-1}, F_d; w) g(m, F_d; \sigma) \right)^h \quad (3)$$

where $C(f_i)$ denotes the children of f_i in all fragmentation graphs containing it, and $C'(f_i) = \{f_i\} \cup C(f_i)$.

However we are unable to maximize this function in closed form. Instead we use the iterative Expectation Maximization (Dempster et al. 1977) technique.

Expectation maximization (EM) In the E-step, the expected log likelihood expression is given by

$$\begin{aligned} Q(w^t, w^{t-1} | \mathcal{X}) &= \mathbb{E}_{w^{t-1}} (\log \mathcal{L}(w^t, \mathcal{X})) \\ &= \sum_{F_1} \dots \sum_{F_d} \Pr(F_1 \dots F_d | \mathcal{X}; w^{t-1}) \log \mathcal{L}(w^t, \mathcal{X}), \end{aligned} \quad (4)$$

where w^t denotes the values for w on the t -th iteration. Substituting (1) and (2) into the above and re-arranging in terms of all possible fragment pairs gives

$$Q(w^t, w^{t-1} | \mathcal{X}) = \sum_{(f_i, f_j) \in \mathcal{F} \times \mathcal{F}} v_{w^{t-1}}(f_i, f_j, \mathcal{X}) \log \rho(f_i, f_j; w^t) + K \quad (5)$$

where

$$\begin{aligned} v_{w^{t-1}}(f_i, f_j, \mathcal{X}) &= \sum_{d'=1}^d \eta_{w^{t-1}}^{d'}(f_i, f_j, \mathcal{X}), \\ \eta_{w^{t-1}}^d(f_i, f_j, \mathcal{X}) &= \sum_{\{(m, h) \in s(x) : x \in \mathcal{X}\}} h \Pr(F_{d-1} = f_i, F_d = f_j | \\ &\quad F_0 = x, P = m; w^{t-1}) \end{aligned}$$

and

$$K = \sum_{F_d} \Pr(F_d | \mathcal{X}; w^{t-1}) \log \Pr(P = m | F_d).$$

In the M-Step, we look for the w^t that maximizes the above expression of Q . Noting that K is independent of w^t and denoting the l th component of w as w_l ,

$$\frac{\partial Q}{\partial w_l} = \sum_{(f_i, f_j) \in \mathcal{F} \times \mathcal{F}} v_{w^{t-1}}(f_i, f_j, \mathcal{X}) \left(\mathbb{I}[f_i \neq f_j] \Phi_{i,j}^l - \sum_{k \in C(f_i)} \Phi_{i,k}^l \rho(f_i, f_k; w) \right) \quad (6)$$

where $\Phi_{i,k}^l$ denotes the l th component of the feature vector $\Phi_{i,k}$ and $\mathbb{I}[\cdot]$ is the indicator function.

This does not permit a simple closed-form solution for w . However $Q(w^t, w^{t-1} | \mathcal{X})$ is concave in w^t , so settings for w^t can be found using gradient ascent. Values for the joint

¹ Although mass spectrometry measures mass over charge, we assume charge is always 1 (see Assumption 1 in Sect. 2.1.1) and hence can use the mass here.

probabilities in the $\eta_{w^{d'-1}}^d$ terms can be computed efficiently using the junction tree algorithm (Koller and Friedman 2009).

We also add an ℓ_2 regularizer on the values of w to Q (excluding the bias term). This has the effect of discouraging overfitting by encouraging the parameters to remain close to zero.

2.2 Combined energy CFM

MS/MS spectra are often collected at multiple collision energies for the same molecule. Increasing the collision energy usually causes more fragmentation events to occur. This means that fragments appearing in the medium and high energy spectra are almost always descendants of those that appear in the low and medium energy spectra, respectively. So the existence of a peak in the medium energy spectrum may help to differentiate between explanations for a related peak in the low or high energy spectra.

For this reason, we also assessed an additional model, combined energy CFM (CE-CFM), which extends the SE-CFM concept by combining information from multiple energies as shown in Fig. 1b. P_{LOW} , P_{MED} and P_{HIGH} each represent a peak from the low, medium and high energy spectrum respectively. The fragment states, transition rules and the observation model are all the same here as for SE-CFM. The main difference now is that the homogeneity assumption is relaxed so that separate transition likelihoods can be learned for each energy block—i.e., F_0 to F_{d_L} , F_{d_L} to F_{d_M} and F_{d_M} to F_{d_H} , where d_L , d_M and d_H denote the fragmentation depths of the low, medium and high energy spectra respectively. This results in separate parameter values for each energy, denoted respectively as w_L , w_M and w_H . The complete parameter set for this model thus becomes $w = w_L \cup w_M \cup w_H$.

We can again use a maximum likelihood approach to parameter estimation based on the EM algorithm. This approach deviates from the SE-CFM method only as follows:

- For each energy level, (6) is computed separately, restricting the $v_{w^{d'-1}}$ terms to relevant parts of the model—e.g. d' would sum from $d_L + 1$ to d_M when computing the gradients for w_M , and from $d_M + 1$ to d_H when computing gradients for w_H .
- The computation of the $\eta_{w^{d'-1}}^d$ terms combines evidence from the full set of three spectra $S(x)$. In SE-CFM, we apply one spectrum at a time, effectively sampling from a distribution over the peaks from each observed spectra. In this extended model we cannot do this because we do not have a full joint distribution over the peaks, but rather we only have marginal distributions corresponding to each spectrum. The standard inference algorithms—e.g. the junction tree algorithm, do

not allow us to deal with observations that are marginal distributions rather than single values. Instead we use the iterative proportional fitting procedure (IPFP) (Deming and Stephan 1940), with minor modifications to better handle cases where the spectra are inconsistent (not simultaneously achievable under any joint distribution). These modifications reassign the target spectra to be the average of those encountered when the algorithm oscillates in such circumstances.

3 Experimental results

In this section we present results using the above described SE-CFM ($d = 2$) and CE-CFM ($d_L = 2$, $d_M = 4$, $d_H = 6$) methods, on a spectrum prediction task, and then in a metabolite identification task.

3.1 Data

We used the Metlin database (Smith et al. 2005), separated into two sets (see description below) each containing positive mode, ESI-MS/MS spectra from a 6510 Q-TOF (Agilent Technologies) mass spectrometer, measured at three different collision energies: 10, 20 and 40 V, which we assign to be low, medium and high energy respectively. Each set was randomly divided into 10 groups for use within a tenfold cross validation framework.

1. *Tripeptides* The Metlin database contains data for over 4,000 enumerated tripeptides. We randomly selected 2,000 of these molecules, then omitted 15 that had four or more rings due to computational resource concerns, leaving 1985 remaining in the set. Fragmentation patterns in peptides are reasonably well understood (Papayannopoulos 1995; Paizs and Suhai 2005), leading to effective algorithms for identifying peptides from their ESI MS/MS data—e.g. (Perkins et al. 1999; Eng et al. 1994; Ma et al. 2003). However, we think that the size of this dataset, and the fact that it contains so many similar yet different molecules, make it an interesting test case for our algorithms.
2. *Metlin metabolites* We use a set of 1,491 non-peptide metabolites from the Metlin database. These are a more diverse set covering a much wider range of molecules. An initial set of 1,500 were selected randomly. Nine were then excluded because they were so much larger than the other molecules (over 1,000 Da), such that their fragmentation graphs could not be computed in a reasonable amount of time.

We also used an additional small validation set, selected because they were measured on a similar mass spectrometer, an Agilent 6520 Q-TOF, but in a different laboratory.




Fig. 4 Two example fragmentations. *a* A non-ring break for which the ion and neutral loss root atoms are labeled. The 1H indicates the movement of a hydrogen to the ion side (marked with +) from the neutral loss side. *b* A ring break for a single aromatic ring of size 6, in which the distance between the broken bonds is 3

These were taken from the MassBank database (Horai et al. 2010). All testing with this set used a model trained for the first cross-fold set of the Metlin metabolite data ($\sim 90\%$ of the data).

3. *MassBank metabolites* This set contains 192 metabolites taken from the Washington State University submission to the MassBank database. All molecules from this submission were included that had MS2 spectra with collision energies 10, 20 and 40 V, in order to provide a good match with the Metlin data. Files containing test molecule lists and assigned cross validation groups are provided as supplementary data.

3.2 Chemical features

The chemical features used in these experiments were as follows. Note that the terms *ion root atom* and *neutral loss (NL) root atom* refer to the atoms connected to the broken bond(s) on the ion and neutral loss sides respectively—cf., Fig. 4.

- *Break atom pair* Indicators for the pair of ion and neutral loss root atoms, each from {C, N, O, P, S, other}, included separately for those in a non-ring break versus those in a ring break—e.g. Fig. 4a: would be non-ring C–C. (72 features)
- *Ion and NL root paths* Indicators for all paths of length 2 and 3 starting at the respective root atoms and stepping away from the break. Each is an ordered double or triple from {C, N, O, P, S, other}, taken separately for rings and non-rings. Two more features indicate no paths of length 2 and 3 respectively—e.g. Fig. 4a): the ion root paths are C–O, C–N and C–N–C. (2020 features).
- *Gasteiger charges* Indicators for the quantised pair of Gasteiger charges (Gasteiger and Marsili 1980) for the

ion and NL root atoms in the original unbroken molecule. (288 features)

- *Hydrogen movement* Indicator for how many hydrogens switched sides of the break and in which direction—i.e. ion to NL (–) or NL to ion (+) {0, ± 1 , ± 2 , ± 3 , ± 4 , other}. (10 features)
- *Ring features* Properties of a broken ring. Aromatic or not? Multiple ring system? Size {3, 4, 5, 6, other}? Distance between the broken bonds {1, 2, 3, 4+}?—e.g. Fig. 4b is a break of a single aromatic ring of size 6 at distance 3. (12 features).

Of these 2,402 features, few take non-zero values for any given break. Many are never encountered in our data set, in which case their corresponding parameters are set immediately to 0. We also append *Quadratic Features*, containing all 2,881,200 pair-wise combinations of the above features, excluding the additional bias term. Again, most are never encountered, so their parameters are set to 0.

3.3 Spectrum prediction

For each cross validation fold, and the MassBank validation set, a model (trained as above), was used to predict a low, medium and high energy spectra for each molecule in the test set. The model is run forward and the resulting marginal distributions for the peak variables are a mixture of Gaussian distributions. We take the means and weights of these Gaussians as our peak mass and intensity values. Since all fragments in the fragmentation graph of a molecule have non-zero probabilities in the marginal distribution, it is necessary to place a cut-off on the intensity values to select only the most likely peaks. Here, we use a post-processing step that removes peaks with low probability, keeping as many of the highest peaks as required to form at least 80 % of the total intensity sum. We also set limits on the number of selected peaks to be at least 5 and at most 30. This ensures that more peaks are included than just the precursor ion, and also prevents spectra occurring that have large numbers of very small peaks. These values were selected arbitrarily, but post-analysis suggests that they are reasonable (see supplementary data). When matching peaks we use a mass tolerance set to the larger of 10 ppm and 0.01 Da (depending on the peak mass), and set the observation parameter σ to be one third of this value. No additional processing was done for the experimental spectra.

3.3.1 Metrics

We consider a peak in the predicted MS/MS spectrum s_P to match a peak in the measured MS/MS spectrum s_M if their masses are within the mass tolerance above. We use the following metrics:

1. *Weighted recall* The percentage of the total peak intensity in the measured spectrum with a matching peak in the predicted spectrum: $100 \times \frac{\sum_{(m,h) \in s_M} h \cdot \mathbb{I}[(m,h) \in s_P]}{\sum_{(m,h) \in s_M} h}$.
2. *Weighted precision* The percentage of the total peak intensity in the predicted spectrum with a matching peak in the measured spectrum: $100 \times \frac{\sum_{(m,h) \in s_P} h \cdot \mathbb{I}[(m,h) \in s_M]}{\sum_{(m,h) \in s_P} h}$.
3. *Recall* The percentage of peaks in the measured spectrum that have a matching peak in the predicted spectrum: $100 \times |s_P \cap s_M| \div |s_M|$.
4. *Precision* The percentage of peaks in the predicted spectrum that have a matching peak in the measured spectrum: $100 \times |s_P \cap s_M| \div |s_P|$.
5. *Jaccard score* $|s_P \cap s_M| \div |s_P \cup s_M|$.

The intensity weighted metrics were included because the unweighted precision and recall values can be misleading in the presence of low-level noise—e.g. when there are many small peaks in the measured spectrum. The weighted metrics place a greater importance on matching higher intensity peaks, and therefore give a better indication of how much of a spectrum has been matched. However, these weighted metrics can also be susceptible to an over-emphasis of just one or two peaks, and in particular of the peak corresponding to the precursor ion. Consequently, we think it is informative to consider both weighted and non-weighted metrics for recall and precision.

3.3.2 Models for comparison

The pre-existing methods,—e.g. MetFrag, FingerID—do not output a predicted spectrum, but skip directly to metabolite identification. So, instead we compare against:

- *Full enumeration* This model considers the predicted spectrum to be one that enumerates all possible fragments in the molecule's fragmentation tree with uniform intensity values.
- *Heuristic (tripeptides only)* This model enumerates known peptide fragmentations as described by (Papayanopoulos 1995), including $b_n, y_n, b_n - H_2O, y_n - H_2O, b_n - NH_3, y_n - NH_3$ and immonium ions.

3.3.3 Results

The results are presented in Fig. 5. For all three data sets, SE-CFM and CE-CFM obtain several orders of magnitude better precision and Jaccard scores than the full

enumerations of possible peaks. There is a corresponding loss of recall. However, if we take into account the intensity of the measured peaks, by considering the weighted recall scores, we see that our methods perform well on the more important, higher intensity peaks. More than 75 % of the total peak intensity in the tripeptide spectra, and ~60 % of the total peak intensity in the metabolite spectra, were predicted.

The results presented in Fig. 5 show scores averaged across the three energy levels for each molecule. If we consider the results for the energy levels separately (see supplementary data), we find that the low and medium energy results are much better for all methods we assessed. For example, in the case of the low energy spectra, the weighted recall scores for SE-CFM are 78, 73 and 81 % for the tripeptide, Metlin metabolite and MassBank metabolite data sets respectively, as compared to 73, 29 and 37 % respectively for the high energy spectra. The poorer high energy spectra results may be due to increased noise and a lower predictability of events at the higher collision energies. Another possible explanation is that the even-electron rule and other assumptions listed in Sect. 2.1.1 may be less reliable when there is more energy in the system. Or perhaps it is simply a factor of the number of peaks per energy level, given that the median numbers of peaks in the measured and predicted spectra respectively were 5 and 6 in the low, 9 and 16 in the medium and 12 and 30 in the high energy spectra.

In the case of the tripeptide data, our methods achieve higher recall scores and similar rates of precision to that of the heuristic model of known fragmentation mechanisms, resulting in improved Jaccard scores. Since peptide fragmentation mechanisms are fairly well understood, this result is not intended to suggest that our method should be used in place of current peptide fragmentation programs, but rather to demonstrate that SE-CFM and CE-CFM are able to extract fragmentation patterns from data to a similar extent to human experts, given a sufficiently large and consistent data set. Like our methods, the heuristic models also perform better for the lower energy levels, with a weighted recall score of 66 % for the low energy, as compared to only 24 % for the high energy.

Unsurprisingly, being a smaller and more diverse data set, the Metlin metabolite results are poorer than those of the tripeptides. However the weighted recall for both our methods is still above 60 % and the precision and Jaccard scores are much higher than for the full enumeration, suggesting that the CFM model is still able to capture some of the common fragmentation trends.

The weighted recall and precision results for the MassBank metabolites are fairly comparable to those of the Metlin metabolites. There is a small loss in the non-




Fig. 5 Spectrum prediction results for tripeptides (*left*), metabolites from Metlin (*middle*) and metabolites from MassBank (*right*). The x-axes show the five metrics: weighted recall (WR), weighted precision (WP), recall (R), precision (P) and Jaccard (J), averaged across the

three energy levels for each test molecule. Bars display mean scores \pm standard error. In each plot, note that the y-axis for Jaccard (on *right*) is different from the others (on *left*)

weighted recall, however this is probably due to a higher incidence of low-level noise in the MassBank data. This results in a small loss in the average Jaccard score. However these results demonstrate that the fragmentation trends learned still apply to a significant degree on data collected at a different time in a different laboratory.

Since this is the first method, to the authors' knowledge, capable of predicting intensity values as well as m/z values, we also investigated the accuracy of CFM's predicted intensity values. We found that the Pearson correlation coefficients for matched pairs of predicted and measured peaks, were 0.7, 0.6 and 0.45 for the low, medium and high spectra respectively (SE-CFM and CE-CFM results were not significantly different). This indicates a positive, though imperfect correlation. Full results and scatter plots are contained in the supplementary data.

Running on a 2.2 GHz Intel Core i7 processor, the median run-time for the spectrum predictions for each molecule in the Metlin metabolite data set was 5 s. Larger molecules with more ring systems generally take longer as they have so many more fragmentation possibilities in the initial enumeration. For molecules with no rings, the median run-time was 2 s, whereas for molecules with 3 or more rings, the median run-time was 9 s. The longest run-time in the Metlin metabolite set was for Troleandomycin (Metlin ID 41012), which has a molecular weight over 800 Da and contains three ring systems, one of which is size 14. It took just under 5 min.

3.4 Metabolite identification

Here we apply our CFM MS/MS spectrum predictions to a metabolite identification task. For each molecule, we produce two candidate sets via queries to two public databases of chemical entities:

1. We query the PubChem compound database (Bolton et al. 2008) for all molecules within 5 ppm of the known molecule mass. This simulates the case where little is known about the candidate compound, but the parent ion mass is known with high accuracy.
2. We query Kyoto Encyclopedia of Genes and Genomes (KEGG) (Kanehisa et al. 2006) for all the molecules within 0.5 Da of the known molecular mass. This simulates the case where the molecule is thought to be a naturally occurring metabolite, but there is more uncertainty in the target mass range.

To conduct this assessment, duplicate candidates were filtered out—i.e. those with the same chemical structure, including those that only differ in their stereochemistry. Charged molecules and ionic compounds were also removed since the program assumes single fragment, neutral candidates (to which it will add a proton). After filtering, the median number of candidates returned from PubChem was 911 for the tripeptides and 1,025 for the metabolites. Note that 9 tripeptides and 57 of the Metlin metabolites were excluded from this testing because no matching entry was

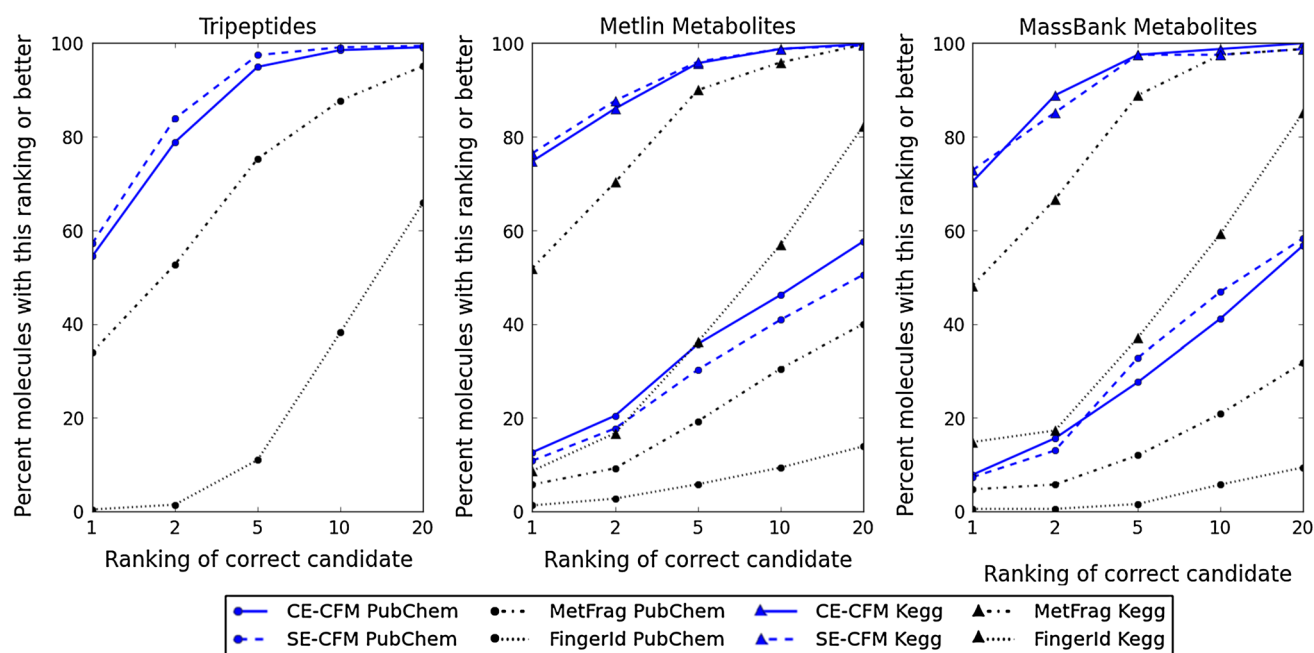


Fig. 6 Ranking results for metabolite identification, comparing both CFM variants with MetFrag and FingerID for tripeptides (*left*), metabolites from Metlin (*middle*) and validation metabolites from MassBank (*right*), querying against PubChem within 5 ppm (*circles*)

and KEGG within 0.5 Da (*triangles*). Note that our methods outperform both MetFrag and FingerID on all metrics, regardless of the database used

found in PubChem for these molecules. The KEGG queries were only carried out for the metabolite data. The median number of candidates returned was 22, however no matching entry was found in KEGG for 833 of the Metlin metabolites and 111 of the MassBank metabolites.

Whenever a matching entry could be found, we ranked the candidates according to how well their predicted low, medium and high spectra matched the measured spectra of the test molecule. The ranking score we used was the Jaccard score described in Sect. 3.3.

We compared the ranking performance of our SE-CFM and CE-CFM methods against those of MetFrag (Wolf et al. 2010) and FingerID (Heinonen et al. 2012). We used the same candidate lists for all programs. For candidate molecules with equal scores, we had each program break ties in a uniformly random manner. This was in contrast to the original MetFrag code, which used the most pessimistic ranking; we did not use that approach as it seemed unnecessarily pessimistic. We set the mass tolerances used by MetFrag when matching peaks to the same as those used in our method (maximum of 0.01 Da and 10 ppm). MetFrag and FingerID only accept one spectrum, so to input the three spectra we first merged them as described by (Wolf et al. 2010): we took the union of all peaks, and then merge together any peaks within 10 ppm or 0.01 Da of one another, retaining the average mass and the maximum intensity of the two. In FingerID we used the linear High Resolution Mass Kernel including both peaks and neutral

losses, and trained using the same cross-fold sets as for our own method. Overall, we attempted to assess CFM, MetFrag and FingerID as fairly as possible, using identical constraints, identical databases and near-identical data input. The results are shown in Fig. 6.

As seen in this figure, our CFM method achieved substantially better rankings than both the existing methods on all three data sets, for both the PubChem and KEGG queries. When querying against KEGG, our methods found the correct metabolite as the top-scoring candidate in over 70 % of cases for both metabolite sets and almost always (> 95 %) ranked the correct candidate in the top 5. In comparison, MetFrag ranked the correct metabolite first in ~50 % of cases for both metabolite sets, and in the top 5 in 89 %. FingerID ranked the correct metabolite first in <15 % of cases.

For PubChem, our methods performed well on the tripeptide data, identifying the correct metabolite as the top-scoring candidate in more than 50 % of cases and ranking the correct candidate in the top 10 for more than 98 % of cases. This is again convincingly better than both MetFrag and FingerID, which rank the correct candidate first in <35 and 2 % of cases respectively.

For the metabolite data, CE-CFM and SE-CFM were able to identify the correct metabolite in only 12 and 10 % of cases respectively, however given that this is from a list of approximately one thousand candidates, this performance is still not bad. Once again, it is substantially better

than MetFrag and FingerID, which correctly identified <6 and 1 % of cases respectively. Our methods rank the correct candidate in the top 10 in more than 40 % of cases on both data sets, as compared to MetFrag's performance of 31 % on the Metlin metabolites and 21 % on the MassBank metabolites. Additionally, the top-ranked compound was found to have the correct molecular formula in more than 88 % of cases for SE-CFM and 90 % of cases for CE-CFM, suggesting that both methods mainly fail to distinguish between isomers. While the performance of all three methods (CFM, MetFrag and FingerID) is not particularly impressive for the PubChem data sets (i.e. <12 % correct) we would argue that the PubChem database is generally a poor database choice for anyone wishing to do MS/MS metabolomic studies. With only 1 % of its molecules having a biological or natural product origin, one is already dealing with a rather significant challenge of how to eliminate a 100:1 excess of false positives. So we would regard the results from the PubChem assessment as a "worst-case" scenario and the results from the KEGG assessment as a more typical metabolomics scenario.

The results for CE-CFM showed minimal difference when compared to those of SE-CFM, casting doubt on whether the additional complexity of CE-CFM is justified. However we think this idea is still interesting as a means for integrating information across energy levels and may yet prove more useful in future work.

The running time of the metabolite identifications is mainly dependent on the number of candidate molecules and the time taken to predict the spectra for each. For example, taking 1,000 candidates (as in the PubChem tests) at the median spectrum prediction run-time of 5 s (see Sect. 3.3), the identification would be expected to take in the order of 1.5 h. Taking only 22 candidates (as in the KEGG tests), this reduces to 2 min. It would be trivial to parallelize the computation by distributing candidates across processors. When repeatedly querying against the same database, it may also be expedient to precompute the predicted spectra to reduce the identification run-time. For example, our web server interface <http://cfmid.wishartlab.com> provides access to precomputed spectra for all 40,000 compounds in HMDB and over 10,000 compounds in KEGG. We encourage readers to make use of this web server, as well as our executables and source code, made available at <http://sourceforge.net/projects/cfm-id>.

4 Conclusion

We have proposed a model for the ESI-MS/MS fragmentation process and a method for training this model from data. The performance has been benchmarked in cross validation testing on a large molecule set, and further

validated using an additional dataset from another laboratory. Head-to-head comparisons using multiple data sets under multiple conditions show that the CFM method significantly outperforms existing state-of-the-art methods, and has attained a level that could be useful to experimentalists performing metabolomics studies.

Acknowledgments Many thanks to Dale Schuurmans, Liang Li, and Jun Peng at the University of Alberta, as well as to the Steinbeck Group at the European Bioinformatics Institute (EMBL-EBI), for invaluable discussions and advice. This work was supported by the Natural Sciences and Engineering Research Council of Canada; Alberta Innovates Technology Futures; and Alberta Innovates Health Solutions and made possible by the Compute Canada Westgrid facility.

References

- Böcker, S., & Rasche, F. (2008). Towards de novo identification of metabolites by analyzing tandem mass spectra. *Bioinformatics*, 24(16), i49–i55.
- Bolton, E., Wang, Y., Thiessen, P., & Bryant, S. (2008). PubChem: Integrated platform of small molecules and biological activities. In Chapter 12 in Annual reports in computational chemistry (Vol. 4). Washington, DC: American Chemical Society.
- Cappé, O., Moulines, E., & Ryden, T. (2005). *Inference in hidden Markov models*. Berlin: Springer.
- de Hoffman, E., & Stroobant, V. (2007). *Mass spectrometry: Principles and applications* (3rd ed.). Chichester: Wiley.
- Deming, S., & Stephan, W. (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *Annals of Mathematical Statistics*, 11(4), 427–444.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B (Methodological)*, 39(1), 1–38.
- Dunn, W. B., & Ellis, D. I. (2005). Metabolomics: Current analytical platforms and methodologies. *Trends in Anal Chem*, 24(4), 285–294.
- Eng, J. K., McCormack, A. L., & Yates, J. R. (1994). An approach to correlate Tandem mass spectral data of peptides with amino acid sequences in a protein database. *Journal of the American Society for Mass Spectrometry*, 5(11), 976–989.
- Fiehn, O. (2002). Metabolomics-the link between genotypes and phenotypes. *Plant Molecular Biology*, 48(1–2), 155–171.
- Galezowska, A., Harrison, M. W., Herniman, J. M., Skylaris, C. K., & Langley, G. J. (2013). A predictive science approach to aid understanding of electrospray ionisation tandem mass spectrometric fragmentation pathways of small molecules using density functional calculations. *Rapid Communications in Mass Spectrometry (RCM)*, 27(9), 964–970.
- Gasteiger, J., & Marsili, M. (1980). Iterative partial equalization of orbital electronegativity: A rapid access to atomic charges. *Tetrahedron*, 36(22), 3219–3228.
- Gasteiger, J., Hanebeck, W., & Schulz, K. P. (1992). Prediction of mass spectra from structural information. *Journal of Chemical Information and Computer Sciences*, 32, 264–271.
- Hastings, J., de Matos, P., & Dekker, A. (2013). The ChEBI reference database and ontology for biologically relevant chemistry: Enhancements for 2013. *Nucleic Acids Research*, 41(Database issue), D456–D463.
- Heinonen, M., Rantanen, A., Mielikainen, T., et al. (2008). FiD: a software for ab initio structural identification of product ions

- from tandem mass spectrometric data. *Rapid Communications in Mass Spectrometry*, 22, 3043–3052.
- Heinonen, M., Shen, H., Zamboni, N., & Rousu, J. (2012). Metabolite identification and molecular fingerprint prediction through machine learning. *Bioinformatics*, 28(18), 2333–2341.
- Hill, A. W., & Mortishire-Smith, R. J. (2005). Automated assignment of high-resolution collisionally activated dissociation mass spectra using a systematic bond disconnection approach. *Rapid Communications in Mass Spectrometry*, 19(21), 3111–3118.
- Horai, H., Arita, M., Kanaya, S., et al. (2010). MassBank: A public repository for sharing mass spectral data for life sciences. *Journal of Mass Spectrometry*, 45(7), 703–714.
- Hufsky, F., Scheubert, K., & Böcker, S. (2014). Computational mass spectrometry for small-molecule fragmentation. *Trends in Analytical Chemistry*, 53, 41–48.
- Kanehisa, M., Goto, S., Hattori, M., et al. (2006). From genomics to chemical genomics: New developments in KEGG. *Nucleic Acids Research*, 34, D354–D357.
- Kangas, L. J., Metz, T. O., Isaac, G., Schrom, B. T., Ginovska-Pangovska, B., Wang, L., et al. (2012). In silico identification software (ISIS): A machine learning approach to tandem mass spectral identification of lipids. *Bioinformatics*, 28(13), 1705–1713.
- Katajamaa, M., & Oresic, M. (2007). Data processing for mass spectrometry-based metabolomics. *Journal of Chromatography A*, 1158(1–2), 318–328.
- Kerber, A., Meringer, M., & Rücker, C. (2006). CASE via MS: Ranking structure candidates by mass spectra. *Croatica Chemica Acta*, 79(3), 449–464.
- Kind, T., & Fiehn, O. (2010). Advances in structure elucidation of small molecules using mass spectrometry. *Bioanalytical Reviews*, 2(1–4), 23–60.
- Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. Cambridge: The MIT Press.
- Levsen, K., Schiebel, H. M., et al. (2007). Even-electron ions: A systematic study of the neutral species lost in the dissociation of quasi-molecular ions. *Journal of Mass Spectrometry (JMS)*, 42, 1024–1044.
- Lindsay, R. K., Buchanan, B. G., Feigenbaum, E. A., & Lederberg, J. (1980). *Applications of artificial intelligence for organic chemistry: The DENDRAL project*. New York: McGraw-Hill Book Company.
- Ma, B., Zhang, K., Hendrie, C., Liang, C., et al. (2003). PEAKS: Powerful software for peptide de novo sequencing by tandem mass spectrometry. *Rapid Communications in Mass Spectrometry*, 17(20), 2337–2342.
- McLafferty, F. W., & Turecek, F. (1993). *Interpretation of Mass Spectra* (4th ed.). Mill Valley: University Science Books.
- Mylonas, R., Mauron, Y., Masselot, A., et al. (2009). X-Rank: A robust algorithm for small molecule identification using tandem mass spectrometry. *Analytical Chemistry*, 81(18), 7604–7610.
- Oberacher, H., Pavlic, M., Libiseller, K., et al. (2009). On the inter-instrument and the inter-laboratory transferability of a tandem mass spectral reference library: 2. Optimization and characterization of the search algorithm. *Journal of Mass Spectrometry (JMS)*, 44(4), 494–502.
- Paizs, B., & Suhai, S. (2005). Fragmentation pathways of protonated peptides. *Mass Spectrometry Reviews*, 24(4), 508–548.
- Papayannopoulos, I. (1995). The interpretation of collision-induced dissociation tandem mass spectra of peptides. *Mass Spectrometry Reviews*, 14(April), 49–73.
- Perkins, D. N., Pappin, D. J. C., Creasy, D. M., & Cottrell, J. S. (1999). Probability-based protein identification by searching sequence databases using mass spectrometry data. *Electrophoresis*, 20, 3551–3567.
- Scheubert, K., Hufsky, F., & Böcker, S. (2013). Computational mass spectrometry for small molecules. *Journal of Cheminformatics*, 5(1), 12.
- Smith, C. A., O'Maille, G., Want, E. J., et al. (2005). METLIN: A metabolite mass spectral database. *Therapeutic Drug Monitoring*, 27(6), 747–751.
- Stein, S. E., & Scott, D. R. (1994). Optimization and testing of mass spectral library search algorithms for compound identification. *Journal of the American Society for Mass Spectrometry*, 5(9), 859–866.
- Sumner, L. W., Amberg, A., Barrett, D., et al. (2007). Proposed minimum reporting standards for chemical analysis. *Metabolomics*, 3, 211–221.
- Tautenhahn, R., Cho, K., Uritboonthai, W., et al. (2012). An accelerated workflow for untargeted metabolomics using the METLIN database. *Nature Biotechnology*, 30(9), 826–828.
- Wishart, D. S. (2011). Advances in metabolite identification. *Bioanalysis*, 3(15), 1769–1782.
- Wishart, D. S., Knox, C., Guo, A. C., et al. (2009). HMDB: A knowledgebase for the human metabolome. *Nucleic Acids Research*, 37, D603–D610.
- Wishart, D. S., Jewison, T., Guo, A. C., et al. (2013). HMDB 3.0: The Human Metabolome Database in 2013. *Nucleic Acids Research*, 41, D801–D807.
- Wolf, S., Schmidt, S., Müller-Hannemann, M., & Neumann, S. (2010). In silico fragmentation for computer assisted identification of metabolite mass spectra. *BMC Bioinformatics*, 11, 148.