

Systems biology

Identification of metabolites from tandem mass spectra with a machine learning approach utilizing structural features

Yuanyue Li ¹, Michael Kuhn^{1,*}, Anne-Claude Gavin^{1,2,†} and Peer Bork^{1,2,3,4,*}

¹Structural and Computational Biology Unit, European Molecular Biology Laboratory, ²Molecular Medicine Partnership Unit (MMPU), 69117 Heidelberg, Germany, ³Max Delbrück Center for Molecular Medicine, 13125 Berlin, Germany and ⁴Department of Bioinformatics, Biocenter, University of Würzburg, 97074 Würzburg, Germany

*To whom correspondence should be addressed.

[†]Present address: Department for Cell Physiology and Metabolism, Centre Médical Universitaire, University of Geneva, Geneva, Switzerland
Associate Editor: Alfonso Valencia

Received on March 20, 2019; revised on July 30, 2019; editorial decision on September 19, 2019; accepted on September 25, 2019

Abstract

Motivation: Untargeted mass spectrometry (MS/MS) is a powerful method for detecting metabolites in biological samples. However, fast and accurate identification of the metabolites' structures from MS/MS spectra is still a great challenge.

Results: We present a new analysis method, called SubFragment-Matching (SF-Matching) that is based on the hypothesis that molecules with similar structural features will exhibit similar fragmentation patterns. We combine information on fragmentation patterns of molecules with shared substructures and then use random forest models to predict whether a given structure can yield a certain fragmentation pattern. These models can then be used to score candidate molecules for a given mass spectrum. For rapid identification, we pre-compute such scores for common biological molecular structure databases. Using benchmarking datasets, we find that our method has similar performance to CSI: FingerID and those very high accuracies can be achieved by combining our method with CSI: FingerID. Rarefaction analysis of the training dataset shows that the performance of our method will increase as more experimental data become available.

Availability and implementation: SF-Matching is available from http://www.bork.embl.de/Docu/sf_matching.

Contact: mkuhn@embl.de or bork@embl.de

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Untargeted mass spectrometry (MS/MS) is a common approach for identification of metabolites in biological samples (Beger *et al.*, 2016; O'Kell *et al.*, 2017; Schrimpe-Rutledge *et al.*, 2016). Thereby, a complex biological sample is analyzed with liquid chromatography electrospray ionization tandem MS/MS, generating several thousands of MS/MS spectra in a few minutes. Inferring all molecular structures from these spectra in a fast, precise manner is, however, still a challenge. The currently fastest way of analyzing such data is to match fragmentation spectra of unknown substances to a reference spectral library (Kind *et al.*, 2018). These spectral libraries are usually built from known purified metabolites or generated by researches experiments. Some databases like METLIN (Guijas *et al.*, 2018), GNPS (Wang *et al.*, 2016) and Massbank (Horai *et al.*, 2010) are collecting these data. This experimental approach has the highest accuracy, however, generating these reference libraries is money- and time-consuming.

Several methods like MS2LDA (van der Hooft *et al.*, 2016, 2017), ChemFrag (Schüler *et al.*, 2018) and ChemDistiller (Laponogov *et al.*, 2018) have been developed for the prediction of functional chemical groups from metabolite spectra. Methods like XCMS2 (Benton *et al.*, 2008) use a similarity-based search to detect possible structural motifs of unknown metabolites in spectra. Also, many methods have been proposed for *in silico* metabolite identification from spectra (Heinonen *et al.*, 2012; Hummel *et al.*, 2010; Nguyen *et al.*, 2018; Vaniya and Fiehn, 2015). *In silico* fragmentation methods, such as ISIS (Kangas *et al.*, 2012), MS-FINDER (Tsugawa *et al.*, 2016), MetFrag (Ruttkies *et al.*, 2016) and CFM-ID (Allen *et al.*, 2015) strive to explain all fragment ions; these methods break every possible covalent bond, scoring each broken bond based on its strength. Some methods like CFM-ID can generate a pre-calculated spectral library. However, as molecular rearrangements occur during fragmentation, precise prediction of the rearrangement is very difficult, and these methods suffer from a low identification accuracy (Blaženović *et al.*, 2018). Other approaches like CSI:

FingerID (Brouard et al., 2016; Dührkop et al., 2015, 2019; Ludwig et al., 2018) convert a spectrum into a fragmentation tree, search this fragmentation tree against a database of known trees, and then infer a molecular fingerprint. These methods need to search fragmentation trees one by one in an online way, which makes it time-consuming when analyzing many spectra.

Depending on the type of biological sample, different chemical search spaces can be used. For well-studied sample types, such as cultured cells or human plasma, many of the metabolites in these samples have been analyzed before. In these cases, spectra can often be matched to reference libraries. Even if reference spectra are not available, there is a high chance that the compounds are part of curated databases such as the KEGG pathway database (Kanehisa et al., 2017) or the HMDB database of metabolites (Wishart et al., 2018). For such compounds, several identification methods have been developed (Blaženović et al., 2018). More complex samples such as those from plants or the environment contain many molecules whose structures have not been determined yet, hampering compound identification using MS.

Here, we describe a new method called SF-Matching (SubFragment-Matching) to predict likely peaks in tandem mass spectra for small molecules using a machine learning approach. Circumventing the complexities of accurately modeling the fragmentation processes and probabilities of bond breakage, our new approach relies on detecting ‘fragile’ substructures in the molecule. These enable us to derive the respective fragmentation patterns to achieve high identification accuracy of compounds from mass spectra.

Unlike similarity-based approaches, which try to annotate spectra with a few features and fail if an unknown metabolite has chemical modifications, SF-Matching can be used for precise identification of a molecule when it contains similar substructures to known metabolites. SF-Matching also does not treat molecular bonds like the *in silico* methods: it does not predict the chemical reaction leading to the fragmentation, it only breaks bonds to find the subfragments, then uses the chemical fingerprint to capture similar molecular structures. Furthermore, the approach described here appears complementary to the existing method CSI: FingerID as a combination with it achieves a much higher accuracy than either method on its own with only a small sensitivity decrease.

2 Materials and methods

2.1 Converting spectra to subfragments

Given a molecular structure, we search for molecular substructures first. We remove all possible bonds $b_{s,t}$ connecting atoms a_s and a_t , with the exception of carbon-carbon bonds in the aliphatic chain and bonds within ring systems. Removing a bond yields two sub-molecules with molecular formulas $F_{s,t}$ and $F_{t,s}$. For each heavy atom a_i in a molecule, the Merck Molecular Force Field (Halgren, 1996) atom type T_{a_i} was determined and the bond type $T_{b_{s,t}}$ was defined as $T_{b_{s,t}} = (T_{a_s}, T_{a_t})$.

The molecule’s corresponding spectrum can be treated as a list of n peaks $\{P_1, P_2, \dots, P_n\}$, here $P_i = (m/z M_i, \text{intensity } I_i)$, and the formula of a fragment ion F_i can be determined from M_i . A mass-to-charge ratio M_i may correspond to several molecular formulas F_i . In this situation, only the formulas that are a subset of the precursor ion are considered. If more than one formula meets this criterion, we include all of them with the same weight.

Then, we test for all $F_{s,t}$ if they are a subset of F_i . In this case, we can calculate the difference of the molecular formulas $\Delta_{i,s,t} = F_i - F_{s,t}$, and define the subfragment $S_j = (T_{b_{s,t}}, \Delta_{i,s,t})$. Therefore, each peak P_i with intensity I_i corresponds to a set of subfragments $\{(S_1, I_i), (S_2, I_i), \dots, (S_n, I_i)\}$.

2.2 Generating models for the subfragments

To build the prediction models, we collected 16 110 molecules for positive ion, 5884 molecules for negative ion from the GNPS, MassBank and in-house databases (Horai et al., 2010; Wang et al., 2016) (Supplementary Tables S1 and S2). For each spectrum in the database, the intensity of the spectrum was normalized so that the

sum over all peak intensities equals one. If one molecule has multiple spectra, they will be treated individually. The spectral peaks were converted to fragment substructures as described earlier. Models were then built for all fragment substructures that occurred in at least five molecules.

Machine learning methods have difficulties within imbalanced training sets. In our case, most molecules do not generate a certain subfragment. To address this imbalance, the training data for a random forest can be augmented with weights for the individual items in the training data. Spectra that contained a peak corresponding to the fragment substructure were assigned the weight equal to the peak’s normalized intensity (I_i). The sum of weights is therefore $w = \sum_i I_i$. For spectra that did not contain a peak corresponding to the fragment substructure were assigned equal weights such that their sum is w . That is, if there are n spectra without the peak, their weight is w/n .

For each molecule, stereoisomer information was removed, and an 8191-bit chemical fingerprint generated using RDKit Fingerprint (Greg Landrum). Then, an extra-trees classifier was built using the scikit-learn (Pedregosa et al., 2012) with 100 trees, using the chemical fingerprints of the complete structure as features and the presence of the fragment substructure as class label. A model will be built if a subfragment existed in at least five molecules. In total, models were built for 1 227 627 subfragments. Each of these models predicts the probability $P(S_{i,x}|C)$ that a peak corresponding to the subfragment $S_{i,x}$ occurs given the chemical fingerprint C .

2.3 Predicting possible peaks for a given molecule

Given a molecule, its chemical fingerprint C was calculated as described earlier. Furthermore, all possible peaks P_i and their corresponding subfragment $S_{i,x}$ were determined from the molecular structure. The molecule’s chemical fingerprint was then used to predict a probability for the existence of a peak, using the pre-built models for the subfragment. When several subfragment had associated models for a given peak, the highest probability of these was assigned to that peak.

$$P(P_i|C) = \max_x P(S_{i,x}|C)$$

2.4 Spectrum scoring

Given a spectrum with peaks P_i normalized peak intensities I_i , the score for a molecule with chemical fingerprint C is calculated by summing over the peak probabilities, using the intensities as weights:

$$\sum_i I_i P(P_i|C)$$

Peaks are determined by searching for molecular formulas within a certain mass accuracy.

2.5 Consensus scoring

For consensus scoring, candidate molecules were scored both with our method as well with CSI: FingerID. A prediction was only accepted if both methods have the same top prediction. When there was a tie for the top prediction, no consensus prediction was recorded.

2.6 Performance evaluation

For the CASMI 2016 dataset, the results of CFM-ID were obtained from the author’s submission, the results of CSI: FingerID were calculated by Sirius 4.0. For the CASMI 2017 dataset, the results of CSI: FingerID were obtained from the author’s submission, the results of CFM-ID were calculated by SE-CFM-trained model. For the EMBL metabolomics core facility (EMBL-MCF) and GNPS datasets, the results of CFM-ID were calculated by SE-CFM-trained model, the results of CSI: FingerID were calculated by Sirius 4.0. If one spectrum had several top predictions, the tie was broken by randomization.

2.7 Data availability and reproducibility

The SF-Matching software and a tutorial can be accessed from http://www.bork.embl.de/Docu/sf_matching. The source code is available at <https://git.embl.de/grp-bork/sf-matching>. An interactive example of using SF-Matching can be accessed at <https://doi.org/10.24433/CO.6279326.v2>. We further deposited the software and pre-calculated spectra predictions for relevant biological molecules [taken from KEGG (Kanehisa *et al.*, 2017), HMDB (Wishart *et al.*, 2018), ChEBI (Hastings *et al.*, 2016) and ChEMBL (Gaulton *et al.*, 2017)] in Zenodo: <http://doi.org/10.5281/zenodo.3345099>.

3 Results

Our concept assumes that molecules consist of fragile and relatively stable substructures. Molecules with similar fragile structures will share similar fragmentation patterns even if they have different stable substructures. For example, in lysophosphatidylinositol and phosphatidylethanolamine, the inositol moiety, ethanolamine and the alkyl chain are stable substructures, connected by a fragile substructure containing ester bonds that are likely to lead to fragmentation. During fragmentation, the two different molecules with different alkyl chains will generate similar fatty acids as fragment ions, although the masses of their fragment ions are different (Fig. 1a). In the fatty acid that can be detected as a fragmentation product, the alkyl chain is the stable substructure and the carboxyl group is the fragment of the fragile substructure.

Our aim was therefore to develop a method that can detect the presence of fragile substructures and based on this predict if a given spectrum is likely to belong to a particular compound. We use machine learning to associate structural information contained in 2D chemical fingerprints with fragmentation patterns. Molecular fingerprints encode the presence of various substructures of the molecule in a vector of bits. If two molecules have similar fragile substructures, some bits of the fingerprint will be the same. As the substructure may vary from molecule to molecule, in our approach, we do not try to identify the extract fragile substructure, but instead use the fingerprint to represent it, and use machine learning to detect the predictive parts of the fingerprint. Given a training database of molecular structures and mass spectra, we process each spectrum individually. On a high level, training the model works as follows (Fig. 1b):

In the molecule associated with the spectrum, we find all substructures by individually breaking all covalent bonds in the molecules. For each broken bond, we record the resulting molecular formulas together with the bond type of the broken bond. For example, $-C_{17}H_{35}$ is one of the substructures of lysophosphatidylinositol.

For all peaks in the spectrum, we calculate the fragmentation products' molecular formulas based on their m/z .

For each fragmentation product, we check which substructures found in Step 1 are a subset of the fragmentation product's molecular formula. For any fragmentation product, there may be several such substructures. Considering each substructure individually, we then designate it as the stable substructure. The remaining part of the fragmentation product must therefore be a part of the fragile substructure, and its molecular formula can be determined from the formulas of the fragmentation product and the stable substructure. Together with the bond type information, we call this part of the fragile substructure a 'subfragment'. Each molecule in the database can be checked if it contains a subfragment.

For each known molecule, we calculate its molecular fingerprint based on the structure of the complete molecule. After determining the presence of subfragments across all training spectra, random forest classifiers are trained for each subfragment based on the molecular fingerprint and the presence of the subfragment in each spectrum.

For testing if a certain molecule is likely to belong to a given spectrum, we do the reverse: we calculate its molecular fingerprint and find all possible subfragments based on the peaks of the spectrum. Then for every subfragment, we use the corresponding random forest classifier to predict the probability of the subfragment.

After finding the possible molecular substructures, we can calculate the mass of all possible fragment ions by adding the formula of subfragment to the formula of substructure.

To increase the speed of the method, we pre-calculated the predicted spectra of all biomolecules in four databases that collect most of the known, relevant biological molecules [KEGG (Kanehisa *et al.*, 2017), HMDB (Wishart *et al.*, 2018), ChEBI (Hastings *et al.*, 2016) and ChEMBL (Gaulton *et al.*, 2017)]. This allows identification of compounds at a rate of more than 10 spectra per second on a laptop with a solid-state drive. The pre-calculated database and searching scripts can be downloaded from http://www.bork.embl.de/Docu/sf_matching.

We evaluated the performance of our approach by comparing it with CFM-ID (Allen *et al.*, 2015), and CSI: FingerID (Dührkop *et al.*, 2015), the two methods that showed very good performance in CASMI 2016 and can also run in batch mode (Schymanski *et al.*, 2017). As many of the test molecules also exist in those software's training dataset, we remove the test spectra from our training dataset to achieve a fair comparison. First, we used all spectra provided in the context of the CASMI 2016 and 2017 automated structural identification challenge (Schymanski *et al.*, 2017) as benchmark dataset. To estimate the performance on multiple chemical databases, we limited the candidates to the molecules that are in the selected database. If one molecule was not present in the target database, the corresponding spectrum was not considered (Supplementary Table S3 and S4). In the CASMI 2016 dataset, SF-Matching had the best performance when searching against four different databases of known molecules (Fig. 2a). In the CASMI 2017 dataset, although the performances of all methods dropped, SF-Matching still showed better performance (Fig. 2b). As an additional benchmark, we also evaluated the methods using spectra from the EMBL-MCF spectral library (Palmer *et al.*, 2018; Supplementary Table S5). We selected candidate molecules from the respective chemical database with an m/z within 5 ppm of target molecules. In this dataset, SF-Matching had a better performance than CSI: FingerID and was also superior to CFM-ID (Fig. 2c).

For this and other approaches, the performance of *in silico* methods drops when removing not only test spectra, but all spectra for these molecules from the training dataset (Dührkop *et al.*, 2019). In CASMI 2016, the paper on CSI: FingerID reported its performance on positive ion when removing part of test molecules from the training dataset (Schymanski *et al.*, 2017). We removed all test molecules from our training dataset and rebuilt the model (Supplementary Table S6). In this dataset, our method has similar performance to CSI: FingerID (Fig. 3a). Based on the published list of CSI: FingerID's training molecules, we also selected spectra from GNPS dataset whose corresponding spectra are not in CSI: FingerID's training dataset (Supplementary Table S7). For this dataset, we also removed molecules from our training dataset and found that our method performed better than CSI: FingerID (Fig. 3b).

As our concept differs from existing ones, we reasoned that it should be possible to achieve a better accuracy if we combine prediction methods. As CSI: FingerID showed good performance in all the three benchmark dataset, we selected spectra where both our method and CSI: FingerID gave the same results. These consensus results achieved about 20% increase in accuracy than any single method, reaching >90% accuracy when analyzing the CASMI 2016 and EMBL-MCF datasets, and still >70% when analyzing the CASMI 2017 dataset. Due to the consensus calculation this comes at the cost of making predictions for fewer spectra, in average, around 55% spectra had consensus identification. The fraction varied between 40 and 80% in the CASMI 2016 and EMBL-MCF datasets and between 10 and 45% in CASMI 2017 (Fig. 4).

SF-matching uses structural features of molecules to build machine learning models and predict the probability that a given molecule spectrum gives rise to a measured spectrum. As machine learning approaches gain power with increasing training sets, we randomly selected subsets of the training dataset to evaluate the performance. Indeed, we observed an increase in the prediction performance with increased training set size (Fig. 5). This result

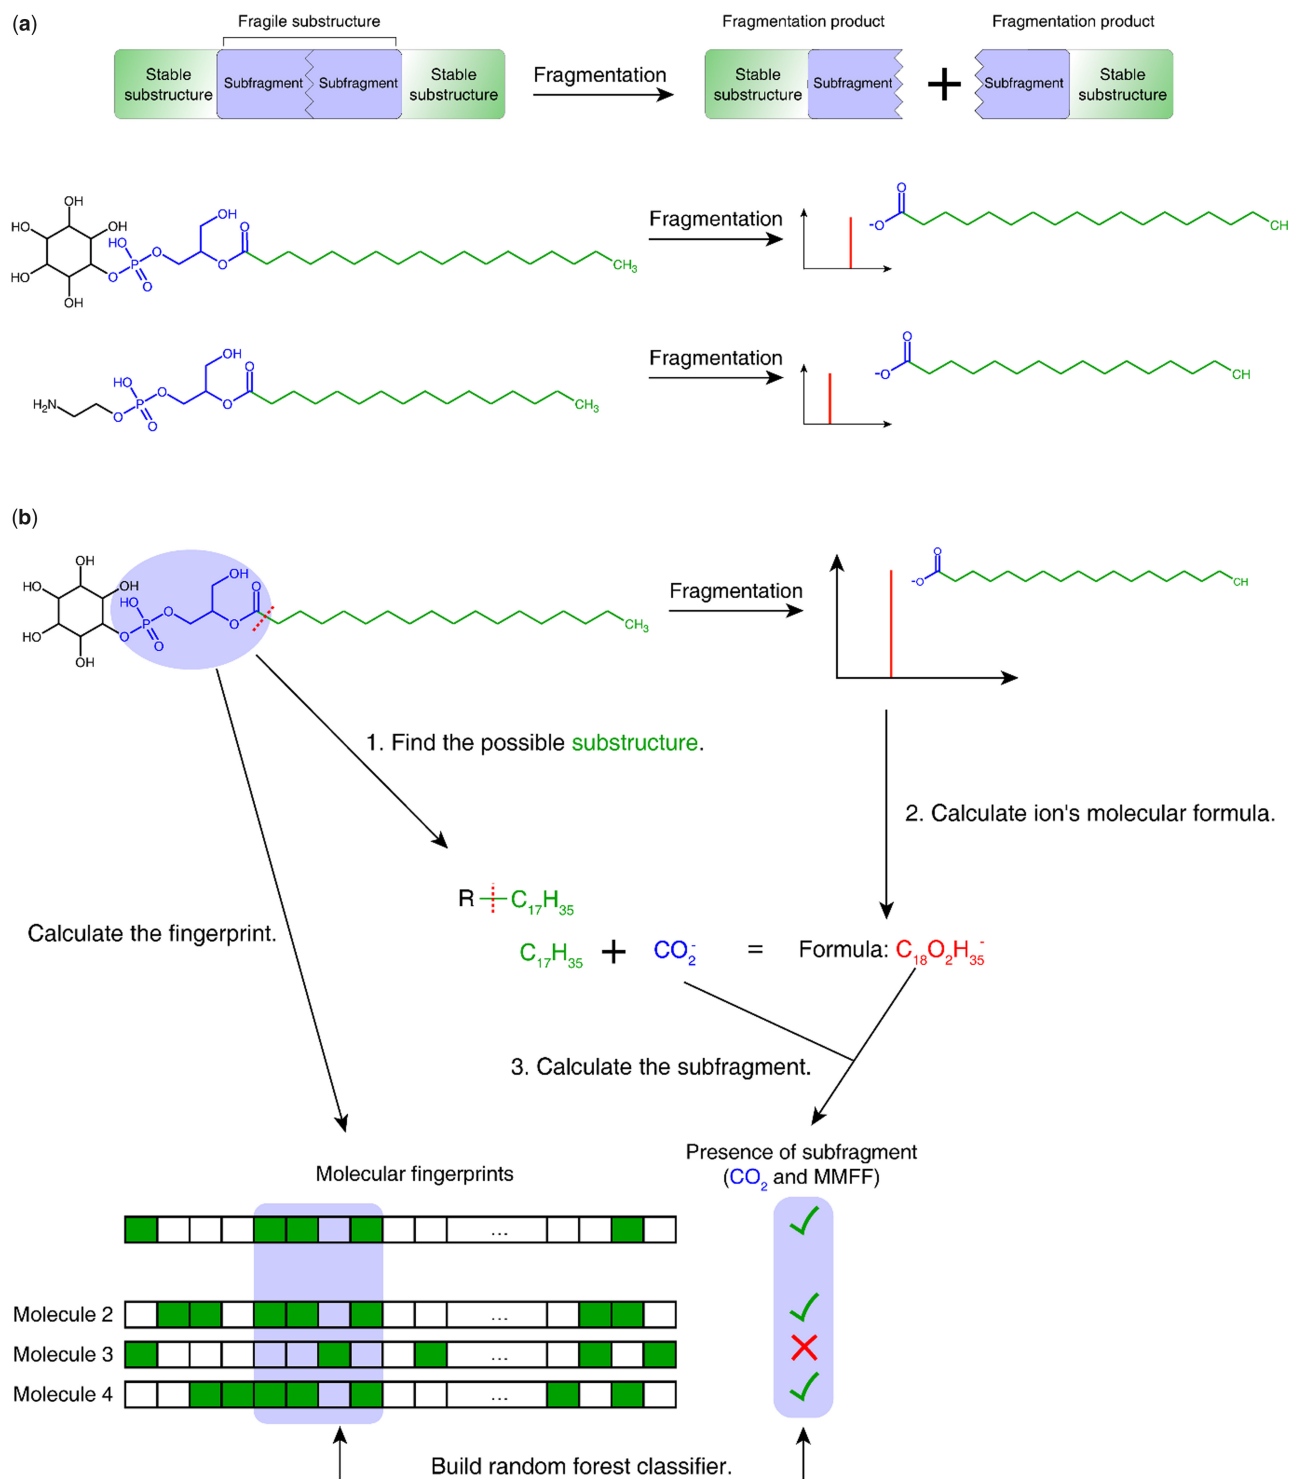


Fig. 1. Schematic of the method. (a) The fragmentation of a molecule can be considered as breaking a fragile substructure surrounded by stable substructures. The resulting fragment contains both a stable substructure and part of the fragile substructure (which we term 'subfragment'). (b) For training the method, we compare possible substructures to peaks in a reference database and calculate the subfragments that were observed. The presence of such a fragile subfragment can then be predicted based on the molecular fingerprint of the complete molecule

suggests that SF-matching will increase performance with time with more experimental spectra becoming available. As shown in [Supplementary Figure S1](#), the training data covers a wide range of the chemical space. Therefore, a general increase in the number of training molecules should be sufficient to increase the performance of SF-Matching.

Taken together, we have developed a new method called SF-Matching to identify the spectra of small molecules in biological samples. Depending on the goals of the MS experiments, SF-matching itself can be used to contribute to candidate molecule identification, given its stand-alone performance, but it can also be used in combination with CSI: FingerID for candidate predictions

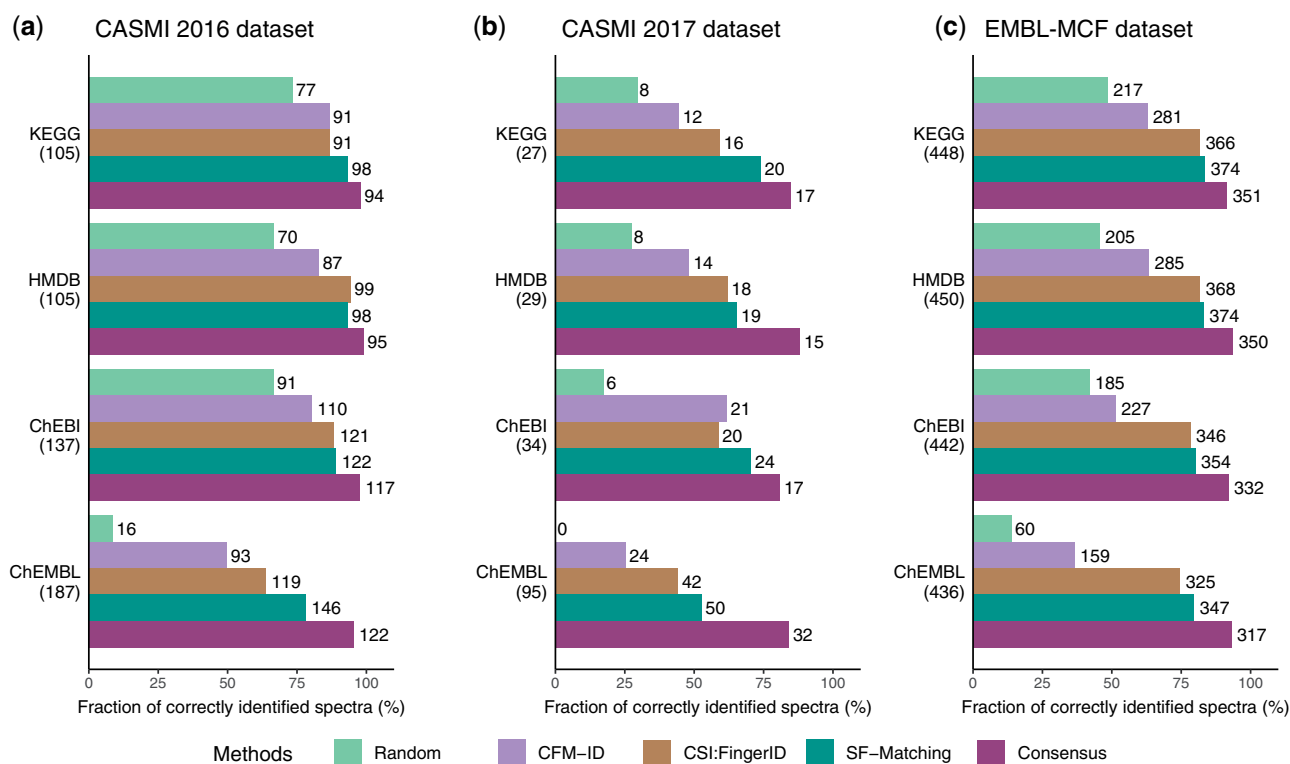


Fig. 2. Performance evaluations when removing the test spectra from training dataset. The accuracy and the sensitivity of SF-Matching, and the consensus method are compared against random prediction and two established methods on (a) the CASMI 2016 dataset, (b) the CASMI 2017 dataset and (c) the EMBL-MCF dataset. The number in the parentheses indicates the total number of molecules which are contained in the various chemical databases; on right of the bar the number of correctly identified molecules is shown. Details on the ranking can be found in Supplementary Tables

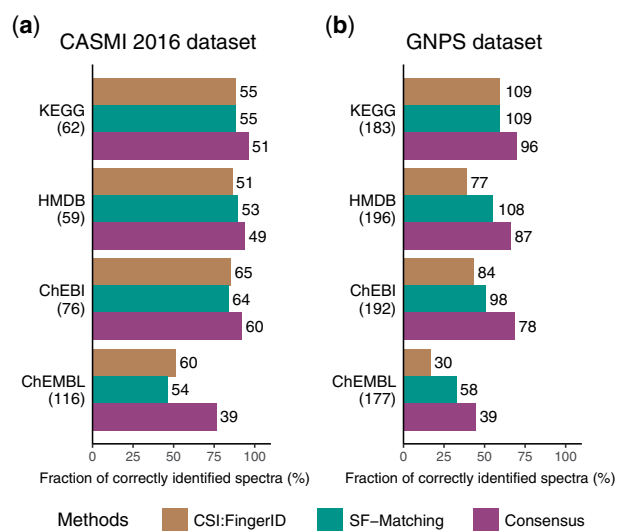


Fig. 3. Performance evaluations when removing the test molecules from training dataset. The accuracy and the sensitivity of SF-Matching, and the consensus method is compared against CSI: FingerID on (a) the positive ion of CASMI 2016 dataset and (b) the GNPS dataset. The number in the parentheses indicates the total number of molecules which are contained in the various chemical databases; on right of the bar the number of correctly identified molecules is shown

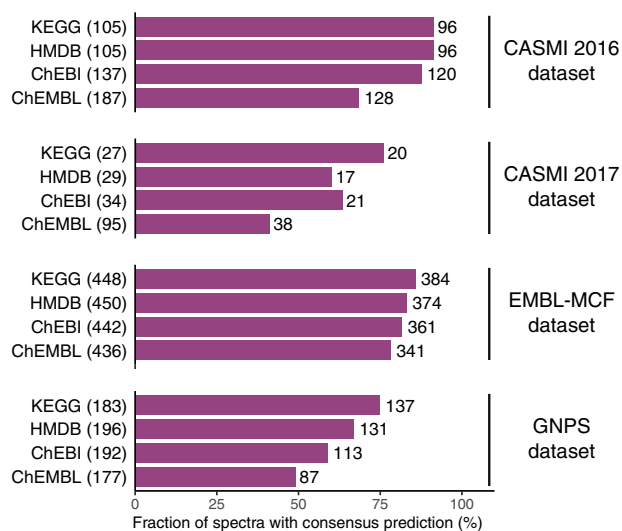


Fig. 4. The fraction of spectra with consensus prediction. The number in the parentheses indicates the total number of molecules which are contained in the various chemical databases; on right of the bar the number of spectra with consensus prediction is shown

Acknowledgements

We thank Joseph Christian Somody, Marco Hennrich, Prasad Phapale and other members of the Bork and Gavin groups as well as the EMBL-MCF for helpful discussions, Enric Mila Vilalta for assistance with the benchmark dataset, and Yan Ping Yuan for IT support.

with high accuracy. Furthermore, as expected from machine learning techniques, the power of the method will increase in the future with the addition of diverse known spectra of biomolecules.

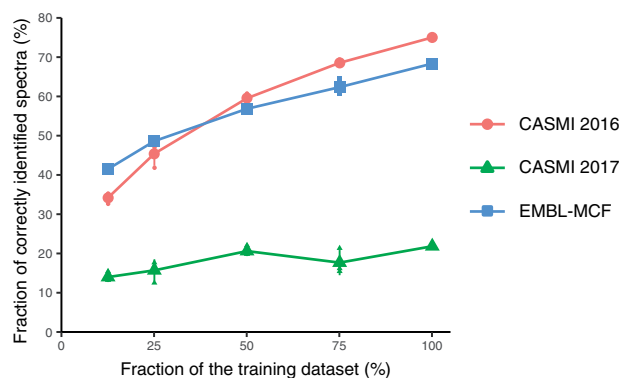


Fig. 5. The effect of training dataset size on the performance on all datasets. For each training dataset size, three training datasets were randomly sampled from the origin whole dataset; here, we use the combination of four databases as candidate database

Funding

We acknowledge funding from EMBL and the MicrobioS grant [ERC-AdG-669830].

Conflict of Interest: none declared.

References

- Allen, F. *et al.* (2015) Competitive fragmentation modeling of ESI-MS/MS spectra for putative metabolite identification. *Metabolomics*, **11**, 98–110.
- Beger, R.D. *et al.* (2016) Metabolomics enables precision medicine: “A White Paper, Community Perspective”. *Metabolomics*, **12**, 149.
- Benton, H.P. *et al.* (2008) XCMS2: processing tandem mass spectrometry data for metabolite identification and structural characterization. *Anal. Chem.*, **80**, 6382–6389.
- Blaženović, I. *et al.* (2018) Software tools and approaches for compound identification of LC-MS/MS data in metabolomics *Metabolites*, **8**, 31.
- Brouard, C. *et al.* (2016) Fast metabolite identification with Input Output Kernel Regression. *Bioinformatics*, **32**, i28–i36.
- Dührkop, K. *et al.* (2015) Searching molecular structure databases with tandem mass spectra using CSI: FingerID. *Proc. Natl. Acad. Sci. USA*, **112**, 12580–12585.
- Dührkop, K. *et al.* (2019) SIRIUS 4: a rapid tool for turning tandem mass spectra into metabolite structure information. *Nat. Methods*, **16**, 299–302.
- Gaulton, A. *et al.* (2017) The ChEMBL database in 2017. *Nucleic Acids Res.*, **45**, D945–D954.
- Guijas, C. *et al.* (2018) Metabolomics activity screening for identifying metabolites that modulate phenotype. *Nat. Biotechnol.*, **36**, 316–320.
- Halgren, T.A. (1996) Merck molecular force field. I. Basis, form, scope, parameterization, and performance of MMFF94. *J. Comput. Chem.*, **17**, 490–519.

- Hastings, J. *et al.* (2016) ChEBI in 2016: improved services and an expanding collection of metabolites. *Nucleic Acids Res.*, **44**, D1214–D1219.
- Heinonen, M. *et al.* (2012) Metabolite identification and molecular fingerprint prediction through machine learning. *Bioinformatics*, **28**, 2333–2341.
- Horai, H. *et al.* (2010) MassBank: a public repository for sharing mass spectral data for life sciences. *J. Mass Spectrom.*, **45**, 703–714.
- Hummel, J. *et al.* (2010) Decision tree supported substructure prediction of metabolites from GC-MS profiles. *Metabolomics*, **6**, 322–333.
- Kanehisa, M. *et al.* (2017) KEGG: new perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Res.*, **45**, D353–D361.
- Kangas, L.J. *et al.* (2012) In silico identification software (ISIS): a machine learning approach to tandem mass spectral identification of lipids. *Bioinformatics*, **28**, 1705–1713.
- Kind, T. *et al.* (2018) Identification of small molecules using accurate mass MS/MS search. *Mass Spectrom. Rev.*, **37**, 513–532.
- Laponogov, I. *et al.* (2018) ChemDistiller: an engine for metabolite annotation in mass spectrometry. *Bioinformatics*, **34**, 2096–2102.
- Ludwig, M. *et al.* (2018) Bayesian networks for mass spectrometric metabolite identification via molecular fingerprints. *Bioinformatics*, **34**, i333–i340.
- Nguyen, D.H. *et al.* (2018) Recent advances and prospects of computational methods for metabolite identification: a review with emphasis on machine learning approaches. *Brief. Bioinform.*, bby066.
- O’Kell, A.L. *et al.* (2017) Untargeted metabolomic analysis in naturally occurring canine diabetes mellitus identifies similarities to human Type 1. *Diabetes. Sci. Rep.*, **7**, 9467.
- Palmer, A. *et al.* (2018) Curatr: a web application for creating, curating and sharing a mass spectral library. *Bioinformatics*, **34**, 1436–1438.
- Pedregosa, F. *et al.* (2012) Scikit-learn: machine Learning in Python. *J. Mach. Learn. Res.*, **12**, 2825–2830.
- RDKit: Open-source cheminformatics; <http://www.rdkit.org>.
- Ruttkies, C. *et al.* (2016) MetFrag relaunched: incorporating strategies beyond in silico fragmentation. *J. Cheminform.*, **8**, 3.
- Schrimpe-Rutledge, A.C. *et al.* (2016) Untargeted metabolomics strategies —challenges and emerging directions. *J. Am. Soc. Mass Spectrom.*, **27**, 1897–1905.
- Schüler, J.A. *et al.* (2018) ChemFrag: chemically meaningful annotation of fragment ion mass spectra. *J. Mass Spectrom.*, **53**, 1104–1115.
- Schymanski, E.L. *et al.* (2017) Critical Assessment of Small Molecule Identification 2016: automated methods. *J. Cheminform.*, **9**, 22.
- Tsugawa, H. *et al.* (2016) Hydrogen rearrangement rules: computational MS/MS fragmentation and structure elucidation using MS-FINDER software. *Anal. Chem.*, **88**, 7946–7958.
- van der Hooft, J.J.J. *et al.* (2016) Topic modeling for untargeted substructure exploration in metabolomics. *Proc. Natl. Acad. Sci. USA*, **113**, 13738–13743.
- van der Hooft, J.J.J. *et al.* (2017) Unsupervised discovery and comparison of structural families across multiple samples in untargeted metabolomics. *Anal. Chem.*, **89**, 7569–7577.
- Vaniya, A. and Fiehn, O. (2015) Using fragmentation trees and mass spectral trees for identifying unknown compounds in metabolomics. *Trends Analyt. Chem.*, **69**, 52–61.
- Wang, M. *et al.* (2016) Sharing and community curation of mass spectrometry data with Global Natural Products Social Molecular Networking. *Nat. Biotechnol.*, **34**, 828–837.
- Wishart, D.S. *et al.* (2018) HMDB 4.0: the human metabolome database for 2018. *Nucleic Acids Res.*, **46**, D608–D617.