MS^2-TRANSFORMER: AN END-TO-END MODEL FOR MS/MS-ASSISTED MOLECULE IDENTIFICATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Mass spectrometry (MS) acts as an important technique for measuring the mass-to-charge ratios of ions and identifying the chemical structures of unknown metabolites. Practically, tandem mass spectrometry (MS/MS), which couples multiple standard MS in series and outputs fine-grained spectrum with fragmental information, has been popularly used. Manually interpreting the MS/MS spectrum into the molecules (i.e., the simplified molecular-input line-entry system, SMILES) is often costly and cumbersome, mainly due to the synthesis and labeling of isotopes and the requirement of expert knowledge. In this work, we regard molecule identification as a spectrum-to-sequence conversion problem and propose an end-to-end model, called MS^2-Transformer, to address this task. The chemical knowledge, defined through a fragmentation tree from the MS/MS spectrum, is incorporated into MS^2-Transformer. Our method achieves state-of-the-art results on two widely used benchmarks in molecule identification. To our best knowledge, MS^2-Transformer is the first machine learning model that can accurately identify the structures (e.g., molecular graph) from experimental MS/MS rather than chemical formula/categories only (e.g., $C_6H_{12}O_6$/organic compound), demonstrating it the great application potential in biomedical studies.

1 INTRODUCTION

Mass spectrometry (MS)-based metabolomics plays a critical role in life science (Klünemann et al., 2021; Banh et al., 2021), which assists in discovering molecular biomarkers/drugs for diagnosis/therapy. MS measures the mass-to-charge (m/z) ratios of adduct ions and outputs a mass spectrum where intensity denotes the ion abundance at the corresponding m/z. In real-world applications, tandem mass spectrometry (MS/MS) is more used, where two standard MS are sequentially coupled for a more fine-grained spectrum than using one MS only (denoted as MS1). In MS/MS, the second MS (denoted as MS2) records the intensities of the fragments which are broken down from the adduct ions in MS1. As a result, MS/MS affords the capability of separating and identifying ions with close m/z-ratios in MS1 based on their fragments in MS2. An example is shown in Figure 1(a), where multiple peaks exist in different m/z ratios. In MS/MS spectrum, the correlations of peaks determine the strengths of chemical bonds, and the locations of peaks indicate the masses of potential functional groups.

To identify the molecules based on MS/MS, conventionally, there are two types of approaches. The first one is accurate mass analysis, in which an experimentally measured mass determines the elemental composition by mass matching. However, such an approach has high false-positive rates, as shown in (Kind & Fiehn, 2006; Shen et al., 2019). The other approach is database searching (DS), in which a matching score between the experimentally obtained spectrum and the database made up of standard spectra is calculated. However, DS does not bring significant improvement in molecule identification due to the limitation of the database. About 90% spectrum of molecules is not covered by the standard database. To overcome this difficulty, researchers turn to more careful chemical designs. For example, people use ^{13}C-labelled isotopes to determine the number of carbons (Tsugawa et al., 2019), but the cost is too high (about $2000 per gram).

Correspondingly, identifying metabolites with machine learning methods is a new trend and progresses have been achieved in metabolite classification by neural networks recently (Dührkop et al., 2021). However, rough classification of chemical categories, like lipids or benzenoids, still cannot

solve the acute problem of identifying the specific structure for profound medical research (e.g., the accurate structure recognition in antibody-antigen interaction), leaving great space for MS/MS identification.

In this work, we regard the task as a spectrum-to-sequence task, where the input is the MS/MS spectrum, and the output is the molecule structure, which is represented by the simplified molecular-input line-entry system (SMILES). Considering that chemists often leverage fragmentation tree (Figure 1b) for molecule identification, which describes the relations of the peaks in the spectrum, we incorporate it into our model. A fragmentation tree can be obtained by the SIRIUS chemical tool (Dührkop et al., 2019).

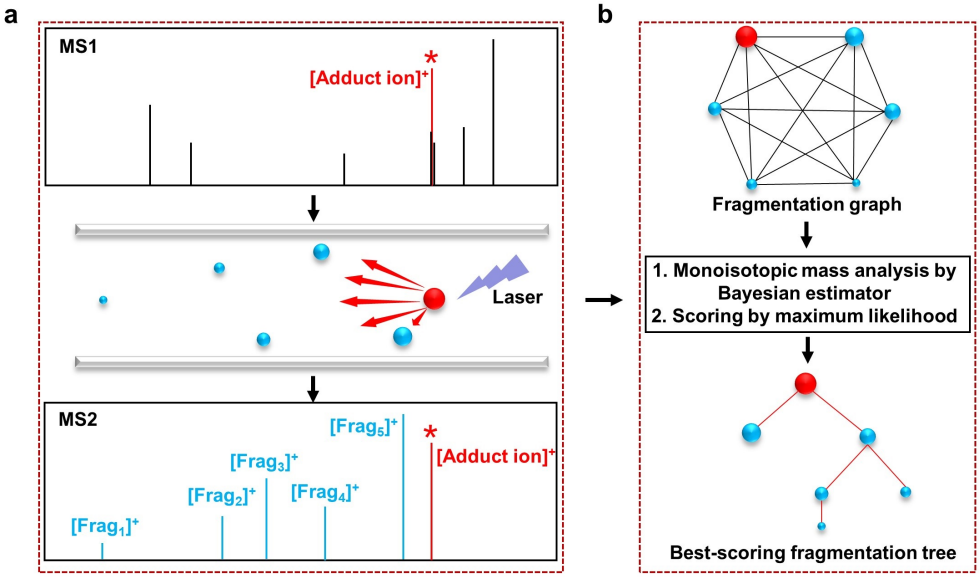

Figure 1: Workflow of MS/MS and fragmentation tree. **(a)** Illustration for adduct ions and fragmentation of metabolites. Metabolites are firstly adducted by selected cations to form [Adduct ion]$^+$ in MS1. Then, [Adduct ion]$^+$ is fragmented to form fragments in an inert gas atmosphere, and the resulted fragments are recorded by MS2. **(b)** Calculation process of fragmentation tree. Based on the MS/MS spectrum, a fully connected graph is initialized for the optimization of the fragmentation tree. Then, a best-scoring fragmentation tree is calculated based on the monoisotopic mass analysis by Bayesian estimator and scoring by maximum likelihood.

We propose an end-to-end model, MS^2-Transformer, for MS/MS assisted molecule identification. MS^2-Transformer follows the encoder-decoder framework, and it takes both the spectrum and the fragmentation tree as input. Our model consists of three modules: (1) *Peak embedding module*, where we map the (m/z ratio, intensity) pairs to high-dimension vectors. The embeddings are used for each peak in the spectrum and each node in the tree. (2) *Fragmentation aggregation module*, where the encoder of MS^2-Transformer consists of stacks of blocks. Each block is made up of a self-attention layer (used to process the spectrum), a graph encoding layer (used to process the fragmentation tree), and a feed-forward layer (used to fuse the features extracted by the above two layers). In this way, the global information in the spectrum and the local information in the tree are fully leveraged. (3) *SMILES reconstruction module*, where the decoder is the standard Transformer decoder, which outputs the SMILES based on the output of the encoder. We conduct experiments on widely used datasets for MS/MS, MassBank and CASMI, and have achieved state-of-the-art results.

The main contributions of our work are as follows:

- We propose MS^2-Transformer, a machine learning-based model to identify the molecule with structural information based on Fragmentation Aggregation (FA) module and Transformer. Experiments on benchmark datasets validate that MS^2-Transformer achieves state-of-the-art performance.

- We incorporate the prior chemical knowledge of MS/MS into the massage passing of local correlations.

- We propose FA module for learning the representation of MS/MS spectrum. The FA performs the massage passing for both local and global correlations on a hypothetical fully connected graph, where nodes denote the peaks in MS/MS.

2 PROBLEM FORMULATION AND NOTATIONS

Let $\mathcal{X}$ and $\mathcal{Y}$ denote source domain and target domain, which are collections of MS/MS spectrum and molecules, respectively. We assume that the molecules are represented by SMILES. Our task is to learn a mapping $f : \mathcal{X} \mapsto \mathcal{Y}$, that can identify the molecule from the input spectrum. We focus on the supervised setting, where the dataset consists of M spectrum-molecule pairs $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{M}$, $x_i \in \mathcal{X}, y_i \in \mathcal{Y}$.

Each x_i can be further decomposed into $x_i = \{(m_j, I_{m_j})\}_{j=1}^{N_i}$, where m_j denotes the m/z ratio, I_{m_j} represents the corresponding ion intensity, and N_i denotes the number of peaks in the spectrum. Note that the m/z ratio is discrete, while the intensity is continuous.

Distinct yet limited patterns can be summarized from MS/MS spectra, providing assistance for the chemists to infer the molecular candidates from bio-fluids. The patterns carrying prior knowledge can be divided into monoisotopic mass and fragmentation patterns.

Monoisotopic mass calculates the molecular mass by summing up the atoms' masses with their most abundant naturally occurring stable isotopes, and it can provide assistance in enumerating all the potential chemical formulas. For example, as shown in Figure 2a, the recorded m/z ratio of 132.1030 Da with a negligible shift of 1.0949 mDa from its monoisotopic mass, could help infer the potential candidate with a chemical formula of $C_6H_{13}NO_2$.

The fragmentation pattern is derived from the correlations among peaks, and it can help find potential functional groups that constitute the molecule. For example, as illustrated in Figure 2a, Δm/z of 18 Da may denote the dehydration process of losing an H_2O, indicating that there exists a hydroxyl in the target molecule's chemical structure.

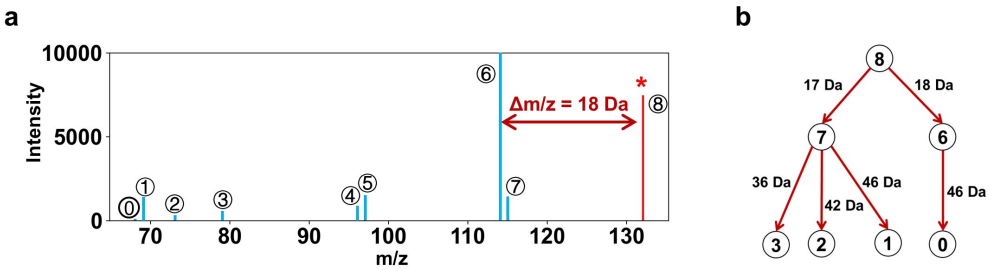

Figure 2: Demonstration of one experimental MS/MS. **(a)** Typical experimental MS/MS of aminocaproic acid, in which the monoisotopic mass is marked with a red star and the fragmentation pattern of dehydration is indicated by Δm/z of 18 Da. **(b)** The fragmentation tree of aminocaproic acid calculated from **(a)**.

Thus, based on the monoisotopic masses and the fragmentation patterns, researchers can enumerate all the possible chemical formulas and functional groups, making it feasible to calculate the hypothetical fragmentation tree by maximum likelihood from the MS/MS spectrum, as shown in Figure 2b. The nodes refer to the peaks in MS/MS, and edges refer to the fragmentation correlations between peaks (Böcker & Dührkop, 2016; Dührkop et al., 2019).

3 OUR MODEL

We introduce our proposed model in this section. Our model consists of three modules, the peak embedding module, the fragmentation aggregation module, and the SMILES reconstruction module. We also introduce a training technique used in our method. The overall framework is in Figure 3.

3.1 PEAK EMBEDDING MODULE

We use an m/z embedding $E_{\text{m/z}} \in \mathbb{R}^{V \times d}$ and an intensity embedding $E_I \times d$ to encode the input. V is the number of distinct m/z ratios. Mathematically, the input m/z m_i (using one-hot representation) and the intensity $I_i \in \mathbb{R}$ is processed as follows:

$$h_i^0 = W_o \texttt{concat}(E_I I_i, E_{\text{m/z}} m_i), \tag{1}$$

where $\texttt{concat}$ means that the inputs are concatenated along the last dimension, $W_o \in \mathbb{R}^{d \times d}$ is the parameter to be learned. The ion intensity I_i is firstly projected into a d-dimension vector, then concatenated with the embedding of m/z, and finally post-processed by a linear transformation parameterized by W_o. In this way, the input x_i is encoded as $(h_1^0, h_2^0, \cdots, h_{N_i}^0)$, where N_i is the number of (m/z ratio, intensity) pairs.

3.2 FRAGMENTATION AGGREGATION MODULE

Given an MS/MS spectrum x_i, the corresponding fragmentation tree T_i can be obtained from x_i with the external chemical tools (Dührkop et al., 2019). Note that the number of nodes in T_i is the same as the number of element in x_i, i.e., N_i. We regard T_i as an undirected graph, and use $\mathcal{E}_i$ to denote the edges in T_i. For each node j in T_i, denote the neighbors as $\mathcal{N}(j) = \{k | (k, j) \in \mathcal{E}_i\}$. The fragmentation aggregation module consists of L blocks, where each block consists of a local aggregation module (i.e., a graph layer), a global aggregation module (i.e., a self-attention layer) and a module to merge the two representations.

Local aggregation: We performed the local information aggregation on the fragmentation tree T_i with MPNN (Gilmer et al., 2017). Let h_i^l denote the features of node i output by the l-th block. We set $\{h_i^0\}$ as the embeddings output by the peak embedding module. Mathematically, the l-th block of local aggregation works as follows:

$$h_{\text{local},i}^l = W_s h_i^{l-1} + \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} W_n h_j^{l-1}, \tag{2}$$

where W_s and W_n are network parameters. That is, we aggregate local features on the tree using GraphSage (Hamilton et al., 2018), and obtain $h_{\text{local},i}^l$.

Global aggregation: In parallel, we use a self-attention layer to globally process all h_i^{l-1}'s. A self-attention layer is defined as follows:

$$h_{\text{global},i}^l = \sum_{j=1}^{N} \alpha_j W_v h_j^{l-1};\ \alpha_j = \frac{\exp(h_i^{l-1} W_q h_j^{l-1})}{\sum_{k=1}^{N} \exp(h_i^{l-1} W_q h_k^{l-1})}, \tag{3}$$

where W_v and W_q are parameters to be learned.

After that, we utilize both local and global features of the l-th block and obtain

$$h_i^{\text{all}} = h^{l-1} + \texttt{LN}(h_{\text{local},i}^l + h_{\text{global},i}^l), \tag{4}$$

where $\texttt{LN}$ denotes layer normalization (Ba et al., 2016).

After that, h_i^{all} is fed into a feed-forward layer and h_i^l is obtained. Mathematically,

$$h_i^l = W_2 \texttt{ReLU}(W_1 h_{\text{all},i}^l + b_1) + b_2, \tag{5}$$

where W_1, W_2, b_1, b_2 are to be learned.

We can eventually obtain $H^L = (h_1^L, h_2^L, \cdots, h_N^L)$ from the last block of the encoder.

3.3 SMILES RECONSTRUCTION BLOCK

We use the standard Transformer decoder (Vaswani et al., 2017) to generate the molecule. The decoder consists of a self-attention layer, encoder-decoder attention layer, and feed-forward layer. The H^L will be used in the encoder-decoder attention, and the SMILES will be generated token by token.

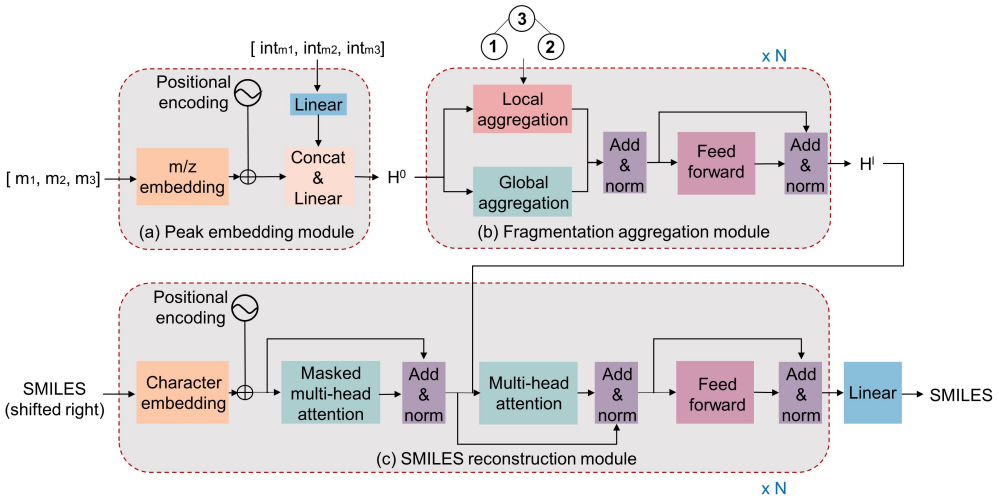

Figure 3: The architecture of MS^2-Transformer. In peak embedding module, the m/z's are processed by an embedding layer with positional encoding, and the intensities are first projected into the same dimension as the m/z's embedding, then concatenated with m/z embedding for the construction of peak embedding by a Linear layer. In the fragmentation aggregation module, the peak embedding was updated by local aggregation on the fragmentation tree, and global aggregation on the fully connected graph, where the nodes denote the peaks in the spectrum and the edges denote the fragmentation correlations. In the SMILES reconstruction module, the learned peak embedding was used for SMILES reconstruction by Transformer.

3.4 DROPPING MODULES

To relieve the overfitting effects, inspired by (Larsson et al., 2017; Zhu et al., 2020), we design an aggregation switch to control the choice of local and global fragmentation aggregation. Mathematically, during training, Eqn.(4) is refined as followed:

$$h_i^{\text{all}} = h_i^{l-1} + \begin{cases} \texttt{LN}(h_{\text{local},i}^l), & 0 \le p < p_{\text{local}}; \\ \texttt{LN}(h_{\text{global},i}^l), & p_{\text{local}} \le p < p_{\text{local}} + p_{\text{global}}; \\ \texttt{LN}(h_{\text{local},i}^l + h_{\text{global},i}^l), & p_{\text{local}} + p_{\text{global}} \le p \le 1; \end{cases} \tag{6}$$

where p is uniformed sampled from $[0, 1]$, $h_{\text{local},i}^l$ and $h_{\text{global},i}^l$ are defined in Eqn.(2) and Eqn.(3) respectively. With probability p_{local}, we only use the local aggregation; with probability p_{global}, we use the global aggregation; otherwise, we use both of them. At inference time, both of the two branches are used as Eqn.(4).

In this work, we focus on spectrum-to-SMILES. The proposed fragmentation aggregation method is general and can be easily applied to other MS/MS-based applications like metabolite classification.

4 EXPERIMENTS

4.1 DATASETS AND BASELINE MODELS

We performed the experiments on the benchmark datasets of electrospray ionization (ESI) MS/MS spectra of the positive ionization mode from MassBank (Horai et al., 2010) and the MS/MS spectra

Table 1: Summary of benchmark datasets.

Dataset	#Pairs	#Edges of T	#Peaks of MS/MS	SMILES length
MassBank	19,729	10.57 ± 12.06	21.82 ± 28.21	35.80 ± 16.75
CASMI	474	11.82 ± 11.78	23.24 ± 27.14	32.27 ± 13.30

of the Critical Assessment of Small Molecule Identification (CASMI) Contest (Nikolic et al., 2017; Schymanski et al., 2017).

The MassBank dataset consists of 19,729 spectrum-SMILES pairs and covers a wide range of chemical compounds, including benzenoids, alkaloids, organic acids, and lipids, etc. The length of SMILES varies from 5 to 100, the number of edges in the fragmentation trees fluctuates from 1 to 59, and the number of peaks in MS/MS covers the range from 2 to 189. The CASMI dataset consists of 474 spectrum-SMILES pairs, and the length of SMILES varies from 5 to 93. The number of edges in the fragmentation trees fluctuates from 1 to 58, the number of peaks in MS/MS covers the range from 2 to 185, and the precursor mass of both MassBank and CASMI is below 500 Da. The detailed statistics can be found in Table 1 and Figure 4.

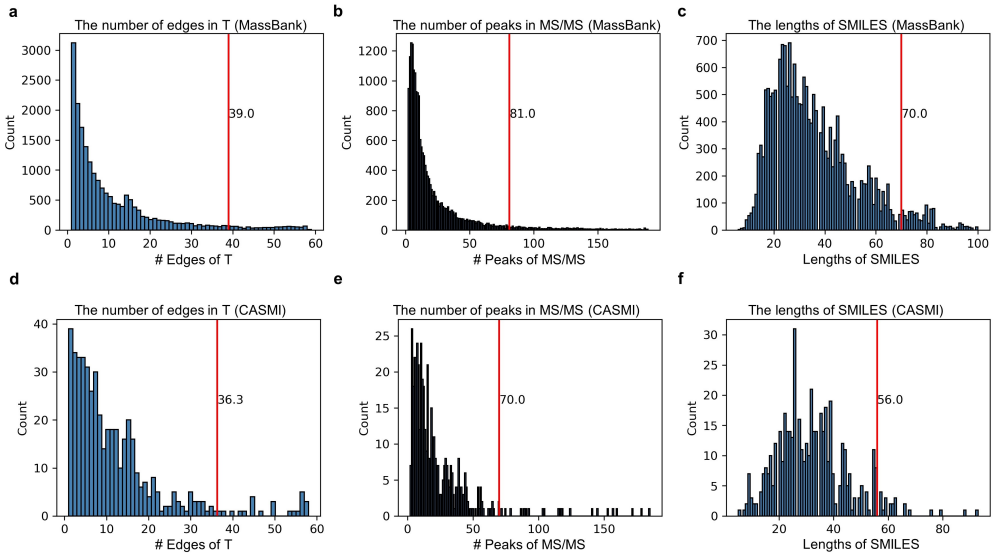

Figure 4: The data statistics of the benchmark datasets. The distribution for the number of **(a)** edges in T and **(b)** peaks in MS/MS, and **(c)** the lengths of SMILES in MassBank dataset. The distribution for the number of **(d)** edges in T and **(e)** peaks in MS/MS, and **(f)** the lengths of SMILES in CASMI dataset. The red line indicated the 95% percentile.

The reconstruction quality of molecule identification was tested based on validity of valency constraints (denoted as Validity), reconstruction accuracy with valid valency (denoted as Reconstruction), and top-1 and -5 reconstruction accuracy (denoted as Top-1 and Top-5). The top-k candidates were generated by beam-search.

We chose the following models as the baselines: **RNN** first encoded the spectrum into hidden representations, then generated the SMILES in a seq2seq architecture with Luong attention mechanism and LSTM operators. We implemented the model using the code shared at `https://github.com/ywk991112/pytorch-chatbot`. **Transformer** represented the classic NLP model, and it first encoded the hidden representations with self-attention mechanism then generated the SMILES with self- and cross-attention mechanisms. We implemented the model using the code shared at `https://github.com/jadore801120/attention-is-all-you-need-pytorch`.

4.2 TRAINING DETAILS

In order to improve generalization and model calibration, we estimate the network's parameters with label smoothing by minimizing the following loss function:

$$y_{ls} = (1-\epsilon) * y_{hot} + \epsilon u(k); \mathcal{L} = -\sum_{k=1}^{K} y_{ls}^{k} \log p_k, \tag{7}$$

where K represented the number of label classes, $u(k)$ represented the prior distribution of labels, and y_{hot} represented the original one-hot encoding.

During training, we used the Adam optimizer (Kingma & Ba, 2017) with varied learning rate (Vaswani et al., 2017) as follows:

$$lr = d^{-0.5} * min(step_num^{-0.5}, step_num * warmup_steps^{-1.5}), \tag{8}$$

where the learning rate was firstly increased during the warm_up steps and then decreased. We trained our models on one machine with four NVIDIA GeForce RTX 2070 Super GPUs with a batch size of 64 and a training epoch of 500.

4.3 EVALUATION RESULTS ON MOLECULE IDENTIFICATION

MassBank: Table 2 shows that MS^2-Transformer has achieved competitive performances on the four metrics with a close number of parameters in the MassBank dataset. As a model introducing the prior chemical knowledge, MS^2-Transformer has achieved comparable performance in the Validity and improved the accuracy by 5% in Reconstruction and top-1 performances compared with Transformer. Our model also achieves a top-5 accuracy of 73.7% for potential chemical candidates.

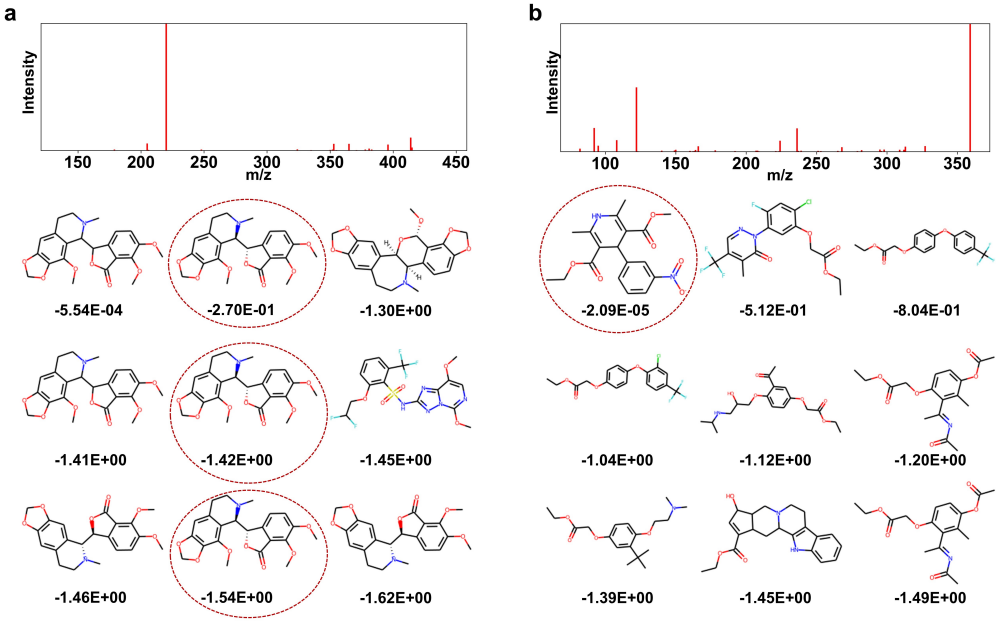

Figure 5: The demonstration of generated SMILES in **(a)** MassBank and **(b)** CASMI dataset. The above spectrum denotes the MS/MS spectrum as the input of MS^2-Transformer, the below molecular graphs denote the top-9 candidates generated with beam search, the number below molecular graph indicates the searching score, and the circled molecule denotes the accurate prediction of true molecule.

CASMI: Table 3 shows the performance of the models in the CASMI dataset. Due to insufficient data pairs in the CASMI dataset, the top-5 identification accuracy only achieves 52.1%. Yet,

Table 2: Comparison of different models' performance on MassBank dataset.

Model	Validity	Reconstruction	Top-1	Top-5
RNN	77.4%	59.0%	45.7%	47.0%
Transformer	**97.7%**	58.0%	56.7%	70.1%
MS^2-Transformer	97.1%	**63.2%**	**61.4%**	**73.7%**

Table 3: Comparison of different models' performance on CASMI dataset.

Model	Validity	Reconstruction	Top-1	Top-5
RNN	78.7%	40.5%	31.9%	31.9%
Transformer	**95.7%**	34.4%	33.0%	46.8%
MS^2-Transformer	**95.7%**	**42.2%**	**40.4%**	**52.1%**

MS^2-Transformer still shows superior performances to Transformer, with an improvement of 8% in Reconstruction performance.

By the way, MS^2-Transformer has obtained the above performances in both MassBank and CASMI only with additional 0.026 M parameters in a light-cost way.

4.4 Fragmentation Inferences

As the correlations of peaks determine the strengths of chemical bonds and indicate the possible functional groups that form the molecule, we also explore MS^2-Transformer's capability in fragmentation inferences.

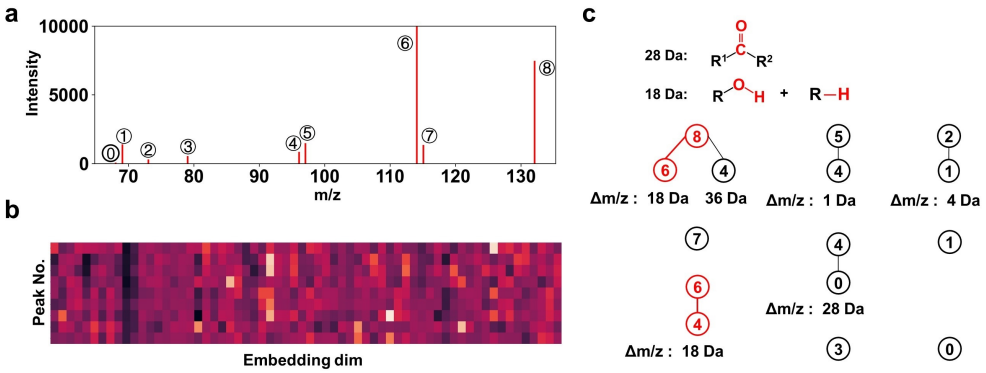

Figure 6: The demonstration of fragmentation inferences. **(a)** The experimental MS/MS spectrum of aminocaproic acid as the input of MS^2-Transformer. **(b)** The peaking embeddings calculated by the fragmentation aggregation module. **(c)** The potential fragmental correlations and possible chemical groups infered from the calculated peak embeddings in **b**.

The fragmentation inference was performed with the following steps: (1) calculate the peak embeddings with the fragmentation aggregation module; (2) calculate the pair-wise cosine similarity matrix of the peak embedding; (3) extract fragmentation correlation with the similarity matrix; (4) annotate the potential chemical fragments.

Figure 6 demonstrates the fragmentation inferences of calculating the potential fragmentation candidates with one single MS/MS. For example, the $\Delta m/z$ of 36 Da between peak 8 and 4 indicates the dehydration reaction of losing two H_2O, which is double validated in the dehydration reaction between peak 8 and 6 of losing one H_2O, and between peak 6 and 4 of losing one H_2O. Figure 7 shows the distribution of fragmentation correlations (represented by the $\Delta m/z$ of parent and child

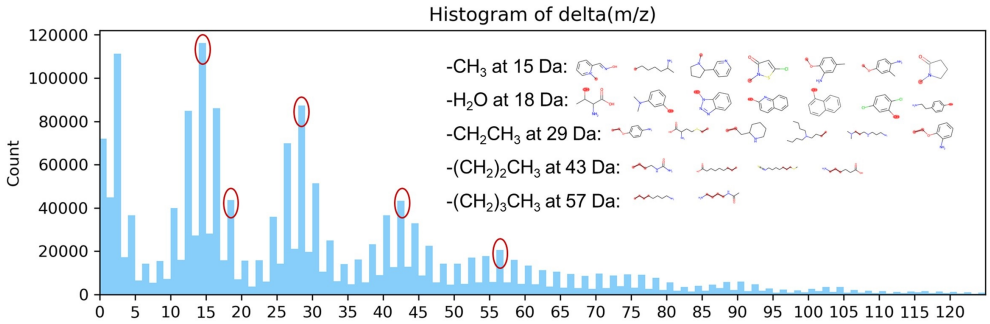

Figure 7: The distribution of the fragmentation correlations in the MassBank dataset infered by MS^2-Transformer. The red circled bars denotes one of the most frequent functional groups, alkyl groups. The molecules with alkyl groups annotated by MS^2-Transformer are listed in the top-right corner.

peaks) in the MassBank dataset. Notably, alkyl affords a higher frequency in the fragmentation correlations. For example, -CH_3 at 15 Da, -CH_2CH_3 at 29 Da, -$(CH_2)_2CH_3$ at 43 Da, and -$(CH_2)_3CH_3$ at 57 Da denote the typical fragmentation patterns of alkyl groups.

5 Conclusions and Future Work

In this paper, we propose the MS^2-Transformer for reconstructing the right chemical structures from MS/MS spectrum based on the aggregation of prior chemical knowledge. Our method identifies the right chemical structures of molecules and exploits the fragmentation inferences for potential fragment candidates. Experimental results showed that MS^2-Transformer outperformed Transformer in the benchmark datasets of molecule identification with an improvement of 5%-7% in top-1 accuracy.

For further improvement, we plan to extend the MS^2-Transformer to problems involving multiple ionization modes (i.e., both positive and negative ionization modes) to investigate the fragmentation aggregation mechanisms in different experimental settings. We also plan to extend our model to encoding the MS/MS spectrum of mixtures, providing help in the clinical research of molecule identification from complex bio-fluids, not only from the pure substances.

References

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016.

Robert S. Banh, Esther S. Kim, Quentin Spillier, Douglas E. Biancur, Keisuke Yamamoto, Albert S. W. Sohn, Guangbin Shi, Drew R. Jones, Alec C. Kimmelman, and Michael E. Pacold. The polar oxy-metabolome reveals the 4-hydroxymandelate coq10 synthesis pathway. *Nature*, 597: 420–425, 2021.

Sebastian Böcker and Kai Dührkop. Fragmentation trees reloaded. *Journal of Cheminformatics*, 8, 2016.

Kai Dührkop, Marcus Ludwig Markus Fleischauer, Alexander A. Aksenov, Alexey V. Melnik, Marvin Meusel, Pieter C. Dorrestein, Juho Rousu, and Sebastian Böcker. SIRIUS 4: a rapid tool for turning tandem mass spectra into metabolite structure information. *Nature Methods*, 16(4): 299–302, 2019.

Kai Dührkop, Louis-Félix Nothias, Markus Fleischauer, Raphael Reher, Marcus Ludwig, Martin A. Hoffmann, Daniel Petras, William H. Gerwick, Juho Rousu, Pieter C. Dorrestein, and Sebastian Böcker. Systematic classification of unknown metabolites using high-resolution fragmentation mass spectra. *Nature Biotechnology*, 39, 2021.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry, 2017.

William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs, 2018.

Hisayuki Horai, Masanori Arita, Shigehiko Kanaya, Yoshito Nihei, Tasuku Ikeda, Kazuhiro Suwa, Yuya Ojima, Kenichi Tanaka, Satoshi Tanaka, Ken Aoshima, Yoshiya Oda, Yuji Kakazu, Miyako Kusano, Takayuki Tohge, Fumio Matsuda, Yuji Sawada, Masami Yokota Hirai, Hiroki Nakanishi, Kazutaka Ikeda, Naoshige Akimoto, Takashi Maoka, Hiroki Takahashi, Takeshi Ara, Nozomu Sakurai, Hideyuki Suzuki, Daisuke Shibata, Steffen Neumann, Takashi Iida, Ken Tanaka, Kimito Funatsu, Fumito Matsuura, Tomoyoshi Soga, Ryo Taguchi, Kazuki Saito, and Takaaki Nishioka. Massbank: a public repository for sharing mass spectral data for life sciences. *Journal of Mass Spectrometry*, 45(7):703–714, 2010.

Tobias Kind and Oliver Fiehn. Metabolomic database annotations via query of elemental compositions: Mass accuracy is insufficient even at less than 1 ppm. *Bioinformatics*, 7, 2006.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

Martina Klünemann, Sergej Andrejev, Sonja Blasche, Andre Mateus, Prasad Phapale, Saravanan Devendran, Johanna Vappiani, Bernd Simon, Timothy A. Scott, Eleni Kafkia, Dimitrios Konstantinidis, Katharina Zirngibl, Eleonora Mastrorilli, Manuel Banzhaf, Marie-Therese Mackmull, Felix Hövelmann, Leo Nesme, Ana Rita Brochado, Lisa Maier, Thomas Bock, Vinita Periwal, Manjeet Kumar, Yongkyu Kim, Melanie Tramontano, Carsten Schultz, Martin Beck, Janosch Hennig, Michael Zimmermann, Daniel C. Sévin, Filipe Cabreiro, Mikhail M. Savitski, Peer Bork, Athanasios Typas, and Kiran R. Patil. Bioaccumulation of therapeutic drugs by human gut bacteria. *Nature*, 597:533–538, 2021.

Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural networks without residuals. In *ICLR*, 2017.

Dejan Nikolic, Martin Jones, Lloyd Sumner, and Warwick Dunn. Casmi 2014: Challenges, solutions and results. *Current Metabolomics*, 5:5–17, 2017.

Emma L. Schymanski, Christoph Ruttkies, Martin Krauss, Céline Brouard, Tobias Kind, Kai Dührkop, Felicity Allen, Arpana Vaniya, Dries Verdegem, Sebastian Böcker, Juho Rousu, Huibin Shen, Hiroshi Tsugawa, Tanvir Sajed, Oliver Fiehn, Bart Ghesquière, and Steffen Neumann. Critical assessment of small molecule identification 2016: automated methods. *Journal of Cheminformatics*, 9, 2017.

Xiaotao Shen, Ruohong Wang, Xin Xiong, Yandong Yin, Yuping Cai, Zaijun Ma, Nan Liu, and Zheng-Jiang Zhu. Metabolic reaction network-based recursive metabolite annotation for untargeted metabolomics. *Nature Communications*, 10, 2019.

Hiroshi Tsugawa, Ryo Nakabayashi, Tetsuya Mori, Yutaka Yamada, Mikiko Takahashi, Amit Rai, Ryosuke Sugiyama, Hiroyuki Yamamoto, Taiki Nakaya, Mami Yamazaki, Rik Kooke, Johanna A. Bac-Molenaar, Nihal Oztolan-Erol, Joost J. B. Keurentjes, Masanori Arita, and Kazuki Saito. A cheminformatics approach to characterize metabolomes in stable-isotope-labeled organisms. *Nature Methods*, 16:295–298, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.

Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tieyan Liu. Incorporating bert into neural machine translation. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=Hyl7ygStwB`.