

• Supplementary File •

ChemDFM-X: Towards Large Multimodal Model for Chemistry

Zihan Zhao^{1†}, Bo Chen^{2†}, Jingpiao Li^{1,2}, Lu Chen^{1,2*}, Liyang Wen², Pengyu Wang^{1,2}, Zichen Zhu¹, Danyang Zhang¹, Yansi Li¹, Zhongyang Dai², Xin Chen^{2*} & Kai Yu^{1,2*}

¹*X-LANCE Lab, Department of Computer Science and Engineering*

MoE Key Lab of Artificial Intelligence, SJTU AI Institute, Shanghai Jiao Tong University, Shanghai 200240, China;

²*Suzhou Laboratory, Suzhou 215123, China*

Appendix A ChemDFM-X

In this section, we detail the methods to develop the ChemDFM-X model. We first give an overview of our training procedure (Section Appendix A.1), introducing what modalities are involved in ChemDFM-X training and how ChemDFM-X is trained. Then, we demonstrate the training details of each modality (Section Appendix A.2~Appendix A.6), including data construction and modality encoder selection.

Appendix A.1 Overview

The overview of the structure and training paradigm of ChemDFM-X is shown in Figure A1. Generally speaking, our ChemDFM-X incorporates the typical “LLM decoder + modality encoder” framework which is widely used by current LMMs [3,19]. Considering the significant differences in data formats among the different modalities, we incorporate separate modality encoders and corresponding projection modules for each modality. Through this “separate encoders + unified decoder” design, the separate encoders enable ChemDFM-X to obtain knowledge and information from different modalities, while the LLM decoder provides the capabilities to aggregate and analyze information from different modalities within a simple generative framework.

Modality Selection. In the field of chemistry, there are primarily two kinds of modalities: *structural* modalities and *characterization* modalities.

Structural modalities mainly directly represent the connections and/or spatial arrangement of molecules and are usually used for reaction inference or theoretical calculation. Among the structural modalities, two modalities are introduced to ChemDFM-X, namely two-dimensional molecular graph and three-dimensional molecular conformation.

On the other hand, characterization modalities mainly imply the partial properties and substructure information of molecules. As the characterization results of molecules, their data is usually in the form of data point sequences with information implicitly hidden among them. One typical usage of characterization modality data is to identify unknown substances. There are many kinds of characterization methods in the field of chemistry. However, due to the expense of chemical experiments, the amount of real experimental data of these modalities is very limited. In this work, we manage to construct a great amount of tandem mass spectrum (MS2) data and infrared spectrum (IR) data, two of the most widely used characterization methods. ChemDFM-X is then trained on these two characterization modalities.

Apart from these four modalities, we also introduce the image modality, including molecular images and reaction images, to ChemDFM-X, as images are the most convenient data used by human researchers. Please refer to Figure ?? for examples of these modalities.

Model Framework. For the LLM decoder, we use one of the advanced chemical LLMs, namely ChemDFM, to leverage the promising chemical language and notation comprehension capability it has acquired. For the modality encoders, according to the current research status of specialist models for different modalities, we either select the advanced existing specialist models or retrain one ourselves as the modality encoders. To better align the output space of modality encoders with the input space of ChemDFM, we incorporate a separate projection module for each modality. Specifically, each of the projection modules is composed of two simple linear layers with a `gelu` activation layer between them.

* Corresponding author (email: chenlusz@sjtu.edu.cn, mail.xinchen@gmail.com, kai.yu@sjtu.edu.cn)

† Zihan Zhao and Bo Chen have the same contribution to this work.

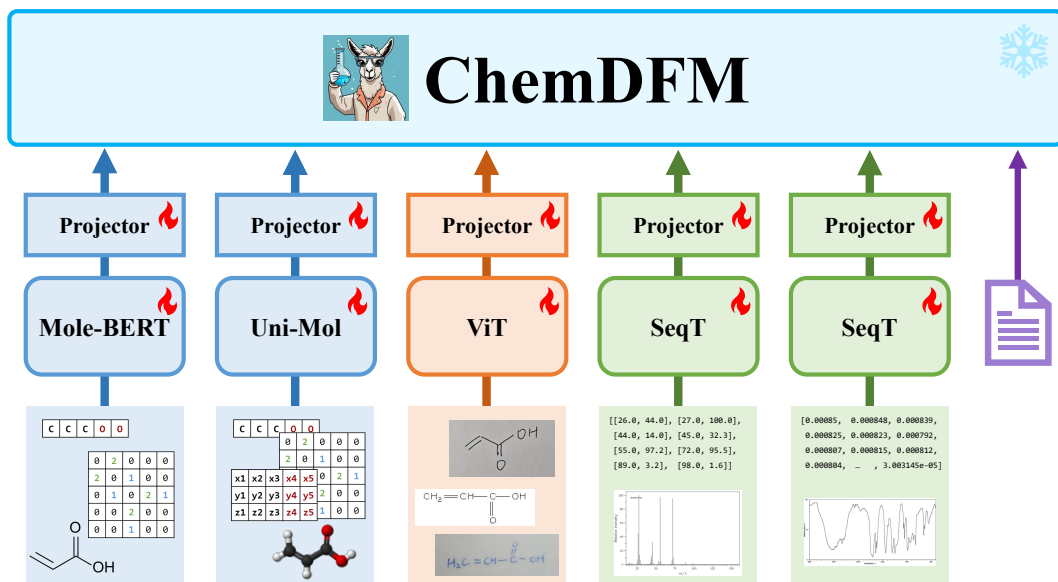


Figure A1 Overview of ChemDFM-X model structure and training paradigm. The colors to mark different input modalities are aligned with Figure 1 in the main text.

Modality	Graph	Conformation	Image	MS2	IR
Modality Encoder	Mole-BERT [38]	Uni-Mol [50]	CLIP [26]	Transformers	Transformers
Projector	2-layer MLP	2-layer MLP	H-Reducer	2-layer MLP	2-layer MLP
Encoding Dims	300	512	1024	768	768
# Modality Tokens	# of Atoms (Dynamic)	# of Atoms (Dynamic)	72 per sub-image (Dynamic)	# of Peaks (Dynamic)	50 (Static)
Peak Learning Rate	1e-5	2e-4	2e-3	2e-4	2e-4
Total Batch Size			256		
Training Epochs			3		

Table A1 Overview of the key hyperparameters of ChemDFM-X.

Training. During the training process of each modal, we freeze the parameters of the pre-trained ChemDFM to maintain the advanced natural language and SMILES processing capabilities it has already acquired. Both the modality encoders and the projection modules are trained to get a better alignment between the encoder outputs and ChemDFM inputs.

The key hyperparameters we chose for each of the modalities are shown in Table A1. The training details of each non-text modality will be introduced in the following subsections, including the instruction-tuning dataset construction process and modality encoder for each modality.

Appendix A.2 Structural Modality: Molecular Graph

The molecular graph is a typical structural modality for molecules. It only consists of the type of atoms in the molecule and the connections between them. Considering that the SMILES notation also contains these two kinds of information, the molecular graph can be easily obtained from the SMILES notation of the corresponding molecule according to simple rules. Formally, a molecular graph can be defined as an undirected graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, where $\mathbf{V} = (v_1, v_2, \dots, v_{|\mathbf{V}|})$ is the atom set with v_i representing the element type of the i -th atom and $\mathbf{E} = (e_1, e_2, \dots, e_{|\mathbf{E}|})$ is the chemical bond set with $e_i = (v_m, v_n, b_i)$ indicating that a chemical bond of type b_i exists between atoms v_m and v_n .

Data Construction. Considering that SMILES notation can also be viewed as a special structural modality and the molecule graphs can be easily obtained from SMILES, the SMILES-related instruction-tuning dataset used by traditional Chemical LLM is naturally suitable for the training of the molecular graph modality. Therefore, we follow the dataset composition and collection procedure proposed by ChemDFM [49] to construct the instruction-tuning data for the molecular graph modality.

Specifically, the data we construct for the instruction tuning of molecular graph modality is mainly composed of the following:

Data Type	# Samples	Data Source	# Heavy Atoms in each Molecule			
			Max	Min	Average	Median
Molecule Description	1152K	PubChem	574	1	35.0	29
Molecular Property Prediction	203K	MoluculeNet	222	1	24.0	22
Reaction Completion	600K	USPTO	273	1	23.3	22
Molecular Graph Recognition	30K	PubChem	150	1	19.5	18

Table A2 Composition of the instruction-tuning data for Molecular Graph and the statistics of the molecules involved

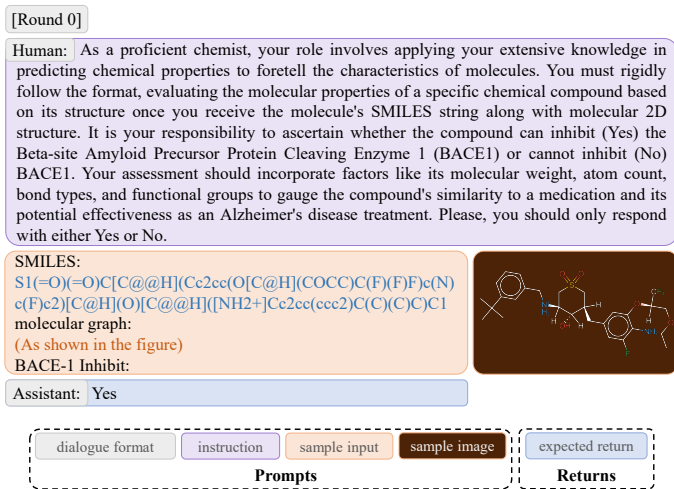


Figure A2 An example of the final structure of instruction tuning data.

- Molecule description.** In this task, models are asked to describe the molecule based on the given molecular representation. The molecule-description pairs are collected from PubChem¹⁾, a web-scale chemical database that contains more than 100M compounds. For each molecule-description pair, we construct two sets of data: one using the molecular graph alone as the molecular representation, and the other using both the molecular graph and the SMILES notation. To further increase the quality of this data, the high-quality molecule-description pairs (where the descriptions have more than two sentences) are repeated twice.

- Molecular property prediction.** In these tasks, models are asked to predict the properties of the given molecule based on its representation. The data are constructed from the most widely used molecular property prediction benchmark, Molecule Net [37]. Similar to the molecule description dataset, we generate the molecular graphs based on the provided SMILES and construct two sets of data using the molecular graph alone or with SMILES.

- Reaction completion.** One of the most important usages of molecular graphs is reaction representation and inference. In this task, models are asked to complete an incomplete reaction where part of the involved molecules are missing. The source reactions are sampled from USPTO [30], the largest chemical reaction database. For each reaction, we randomly mask the reactant, reagent, or product, where the remaining molecules of the reaction are represented by molecular graph alone or with SMILES.

- Molecular graph recognition.** In these tasks, models are asked to recognize the molecule graph by providing the alternative representation of the corresponding molecule, such as SMILES, IUPAC²⁾ name, or molecular formula. The involved molecules are sampled from PubChem. Based on the molecules, we construct five sub-tasks: predicting the IUPAC name given the molecular graph alone or with SMILES, predicting the molecular formula given the molecular graph alone or with SMILES, and predicting the SMILES notation given the molecular graph alone.

It is worth noticing that we carefully remove all the data that may be present in the evaluations based on SMILES matching. For the prompt, we leverage the same dialogue format used by the ChemDFM, while rewriting the instruction of each task using GPT-4 to diversify the expressions. A detailed example of the data in our instruction-tuning dataset is illustrated in Figure A2.

Graph Encoder. We utilize one of the most recent and advanced models, namely Mole-BERT [38], which specializes in 2D molecular graph understanding. The model is an encoder-decoder based graphic neural network (GNN) pre-trained on 2 million molecules and achieves promising performances on multiple downstream tasks. Specifically, the Mole-BERT model follows an encoder-decoder architecture where the encoder and the decoder each consist of 5-layer graph isomorphism

1) <https://pubchem.ncbi.nlm.nih.gov/>

2) short for International Union of Pure and Applied Chemistry

Data Type	# Samples	Data Source	# Heavy Atoms in each Molecule			
			Max	Min	Average	Median
Molecule Description	1152K	PubChem	574	1	35.0	29
Molecular Property Prediction	203K	MoluculeNet	222	1	24.0	22
Reaction Completion	600K	USPTO	273	1	23.3	22
Molecular Conformation Recognition	30K	PubChem	150	1	19.5	18

Table A3 Composition of the instruction-tuning data for Molecular Conformation and the statistics of the molecules involved

Data Type	# Samples	Data Source	# Heavy Atoms in each Molecule			
			Max	Min	Average	Median
Molecule Description	1152K	PubChem	574	1	35.0	29
Molecular Property Prediction	204K	MoluculeNet	222	1	24.0	22
Reaction Completion	600K	USPTO	273	1	23.3	22
Molecular Image Recognition	30K	PubChem	150	1	19.5	18
Reaction Image Recognition	300K	USPTO	440	1	23.7	22

Table A4 Composition of the instruction-tuning data for Image modality and the statistics of the molecules involved

networks (GINs). We adopt its encoder as our molecular graph encoder and utilize the resulting node representations for each atom as the feature of the molecular graphs.

Appendix A.3 Structural Modality: Molecular Conformation

The molecular conformation is another typical structural modality for molecules. Compared with molecular graphs, molecular conformations additionally contain the coordinate information of each atom of the molecules, reflecting the position information of atoms in the molecule. Therefore, molecular conformations usually contain more information than molecular graphs and SMILES notations. Formally, a molecular conformation can be viewed as a similar undirected graph to molecular graphs $\mathcal{G} = (\mathbf{V}', \mathbf{E})$ with $\mathbf{V}' = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{|\mathbf{V}'|})$ and $\mathbf{v}_i = (x_i, y_i, z_i, a_i)$ where (x_i, y_i, z_i) is the Cartesian coordinates of the i -th atom and a_i denotes the element type of the i -th atom.

Data Construction. We construct the instruction-tuning dataset for molecular conformation modality based on the dataset used for molecular graph modality. Specifically, we utilize the same source data as molecular graphs while replacing all the molecular graphs in the inputs with molecular conformations. We also use the same input format while instruction is rewritten specifically for molecular conformations.

However, it is worth noticing that obtaining molecular conformations from SMILES is a lot more complicated than obtaining molecular graphs from SMILES. Ideally, rigorous calculations based on quantum chemistry are needed to acquire the “optimal” conformations where the potential energy of the molecule is the lowest. However, quantum chemistry calculation can be very costly, especially for molecules with a large amount of atoms. To tackle this problem, we follow the same approximate calculation process used by Uni-Mol [50] to obtain the approximated conformations under optimization algorithms. Specifically, we first generate the original molecular conformations using RDKit [17]. Then, the molecular conformation is optimized using the Merck Molecular Force Field (MMFF) [13] algorithm and the pseudo-optimal conformation for the molecule is obtained.

Conformation Encoder. We adopt one of the famous large-scale molecular conformation representation models, Uni-Mol [50], as our 3D molecular conformation encoder. It is a transformer-based 3D molecular pre-trained model derived from 19M molecules and 209M molecular conformations. It achieved promising results in various downstream tasks, especially in 3D spatial tasks.

Appendix A.4 Image Modality

During the work of human chemists, images, rather than structured molecular graphs and conformations, are the most used modality through which they perceive and express molecules and reactions. Therefore, it is vital for the chemical LLM to possess image comprehension capability so that the chemist can communicate with them more conveniently.

Data Construction. We utilize the same data sources as molecular graphs and molecular conformations to construct the instruction-tuning dataset for the image modality by replacing the non-text molecular representation with images and generating corresponding instructions.




Figure A3 Examples of three molecular styles.

Data Type	# Samples	Data Source	# Heavy Atoms in each Molecule			
			Max	Min	Average	Median
MS2 spectrum recognition	21K	PubChem	99	1	18.3	18
Molecular Property Prediction	90K	MoluculeNet	135	1	23.4	22
Reaction Related MS2 Identification	224K	USPTO	131	1	23.5	22

Table A5 Composition of the instruction-tuning data for MS2 Spectrum and the statistics of the molecules involved

Moreover, to diversify the styles of images, we employed three different methods for the generation of molecular images. Firstly, we directly use two different toolkits namely RDKit [17] and Indigo [10]. We further conduct traditional image augmentation methods, such as random grayscale conversion and random color jitter, on the resulting images. Besides, to familiarize ChemDFM-X with the handwritten style images, we utilize the pipeline proposed in [1] to generate pseudo-handwritten images. The comparison of the three styles of images is illustrated in Figure A3.

In addition to the existing tasks that revolve around molecule images, we add a new task focusing on the reaction images. Specifically, we randomly sample 300K reactions from the USPTO database and generate the images of the reactions using the RDKit toolkit. Based on the reaction image alone, the model is asked to recognize the reaction by providing the SMILES notation of the whole reaction.

Image encoder. There is no previous work that specifically targeted both molecule images and reaction images. Therefore, we employ CLIP [26], the image encoder widely used in general domain LMMs. The original CLIP model can only process images with resolutions of 224×224 or 336×336 . It is far from being sufficient for chemical images, especially for images of large molecules and reactions. Pre-experiments show that forcibly compressing images to these resolutions results in the model learning almost nothing. To handle this problem, we follow the method proposed by [19]. Specifically, instead of directly resizing the images of original resolution, we crop them into multiple sub-images that fit the input size of CLIP and then leverage a specific module named H-Reducer [15] to reduce the output sequence length of CLIP by a factor of n . In our cases, n is set to 8.

Appendix A.5 Characterization Modality: MS2 Spectrum

Molecular graphs, molecular conformations, and molecular images provide information about the structure of known molecules, however, numerous molecules of unknown structure exist in nature or may appear during experiments. In this context, spectroscopy becomes one of the key tools to determine the identity of a molecule. Mass spectrum is one of the common types of spectra, which is based on the principle that the molecule is ionized by an ion source and then cleaved into ion fragments with different mass-to-charge ratios.³⁾ The mass spectrum consists of a series of ion peaks, where the horizontal axis represents the mass-to-charge ratio of the ion fragments and the vertical axis represents the intensity of the peaks. Formally, the MS2 spectrum data can be represented by a tuple sequence $\mathbf{M} = ((r_1, I_1), (r_2, I_2), \dots, (r_n, I_n))$, where (r_i, I_i) denotes the i -th peak, r_i denotes the mass-to-charge ratio of the i -th peak, I_i denotes the intensity of the i -th peak, and n represents the total number of peaks in the MS2 spectrum. By introducing mass spectrum data, the ability to deal with unknown molecules can be added to the chemical LLM.

Data Construction. As a typical kind of characterization modalities, the tasks suitable for MS2 spectra differ from those for the structural modalities. Therefore, to construct the instruction-tuning dataset for MS2 spectra, we select a small subset of tasks that are suitable for MS2 spectra from our instruction-tuning dataset for structural modalities, and

3) https://en.wikipedia.org/wiki/Mass_spectrum

Data Type	# Samples	Data Source	# Heavy Atoms in each Molecule			
			Max	Min	Average	Median
IR spectrum recognition	30K	PubChem	150	1	19.5	18
Molecular Property Prediction	102K	MoluculeNet	222	1	24.0	22
Reaction Related IR Identification	300K	USPTO	440	1	22.9	22

Table A6 Composition of the instruction-tuning data for IR Spectrum and the statistics of the molecules involved

supplement them with tasks specific to characterization modalities. Specifically, the instruction-tuning dataset for MS2 spectra is composed of the following tasks:

- **MS2 spectrum recognition.** We utilize the same data source as the molecular graph recognition task introduced in Section Appendix A.2 with the molecular graph replaced by the MS2 spectra of the corresponding molecule.

- **Molecular property prediction.** The constructing process is also akin to that of the molecular graph modality. However, different from molecular graphs, it does not make sense to predict the molecular properties using MS2 spectra alone. Therefore, we only instruct the model to conduct the prediction based on the MS2 spectra along with the SMILES notations.

- **Reaction related MS2 spectrum identification.** In practice, one of the most important usages of the spectra is to identify unknown molecules in the context of a reaction. In such scenarios, the identity of the unknown molecules should be inferred in accordance with not only the spectra, but also the a priori information of the reaction, like the known reactants or the detected products. Therefore, the capability to identify the MS2 spectrum with the help of the reaction related to the molecule is also important and valuable for chemical LMMs. To construct the data for such tasks, we reuse the data for reaction completion tasks in the instruction-tuning datasets of structural modalities. The incomplete reactions will be represented by SMILES notations, while the MS2 spectra of the missing substances are also provided to the model. Based on this information, the model is instructed to identify the corresponding molecules by providing the SMILES of the molecules.

Another great challenge for the MS2-spectrum modality is the great expense associated with obtaining real experimental data. That is because of not only the high cost of MS2 spectroscopy, but also the significant human labor required throughout the experimental process. Therefore, large-scale experimental MS2 spectra data is absent. To tackle this problem, we use a prediction-based approach to generate mass spectrum data in large batches. The prediction tool we used is CFM-ID 4.0 [35], which provides a method for accurate and efficient prediction of molecular MS2 spectra by creating a probabilistic graphical generative model for the MS2 fragmentation process through competitive fragmentation modeling, and adjusting the model parameters from the data through machine learning algorithms. With CFM-ID 4.0, we obtain a total of about 700K tandem mass spectra with their corresponding SMILES to form the instruction tuning dataset for for training the encoder of MS2 spectrum modality .

MS2 Sequence Transformer. Since there is little work on encoding MS2 spectra, we choose to train a new spectrum encoder for MS2. Considering the data format of MS2 spectra is point sequences and one of the classical applications of the spectrum is structural inference [8], we train an encoder-decoder manner transformer model with the pertaining task designed as a spectral inference task. Specifically, the model is trained to generate the SMILES notation of the molecule corresponding to the input MS2 spectrum. We build the codebook of the tokens based on the transverse coordinates and then acquire the input token list from the ion peaks of the MS2 spectrum using the codebook. After the model is trained, we use its encoder module as the encoder for the MS2 spectrum modality.

Appendix A.6 Characterization Modality: IR Spectrum

Similar to the MS2 spectrum, the IR spectrum also contains rich chemical information about the molecule and is one of the most important means to characterize the molecular structure. The principle of the infrared spectrum is that a molecule can selectively absorb infrared rays of certain wavelengths, causing its own vibrational and rotational energy levels transitions, which is related to the detailed structure of molecules⁴). The infrared spectra of molecules can be obtained through the detection of infrared absorption. The horizontal coordinate of the infrared spectrum indicates the wavelength of the infrared ray, and the vertical coordinate generally implies the absorption intensity. Formally, the IR spectrum can be represented by a point sequence $\mathbf{R} = ((w_1, t_1), (w_2, t_2), \dots, (w_l, t_l))$, where w_i denotes the wave length of the light, t_i is the absorption intensity of the corresponding light.

Data Construction. We construct the instruction-tuning dataset for the IR-spectrum modality based on the dataset used for the MS2-spectrum modality. Specifically, we utilize the same source data while replacing all the MS2 spectra with IR spectra with corresponding modifications of the instructions.

Similar to MS2 spectra, IR spectra also suffer from the high cost of the experimental data. To make things even worse, the expense of IR spectra calculation is also beyond acceptable. Because the simulation of molecular infrared spectra needs to go through the time-consuming DFT process and the large-volume infrared computation puts higher demands on the

4) https://en.wikipedia.org/wiki/Infrared_spectroscopy

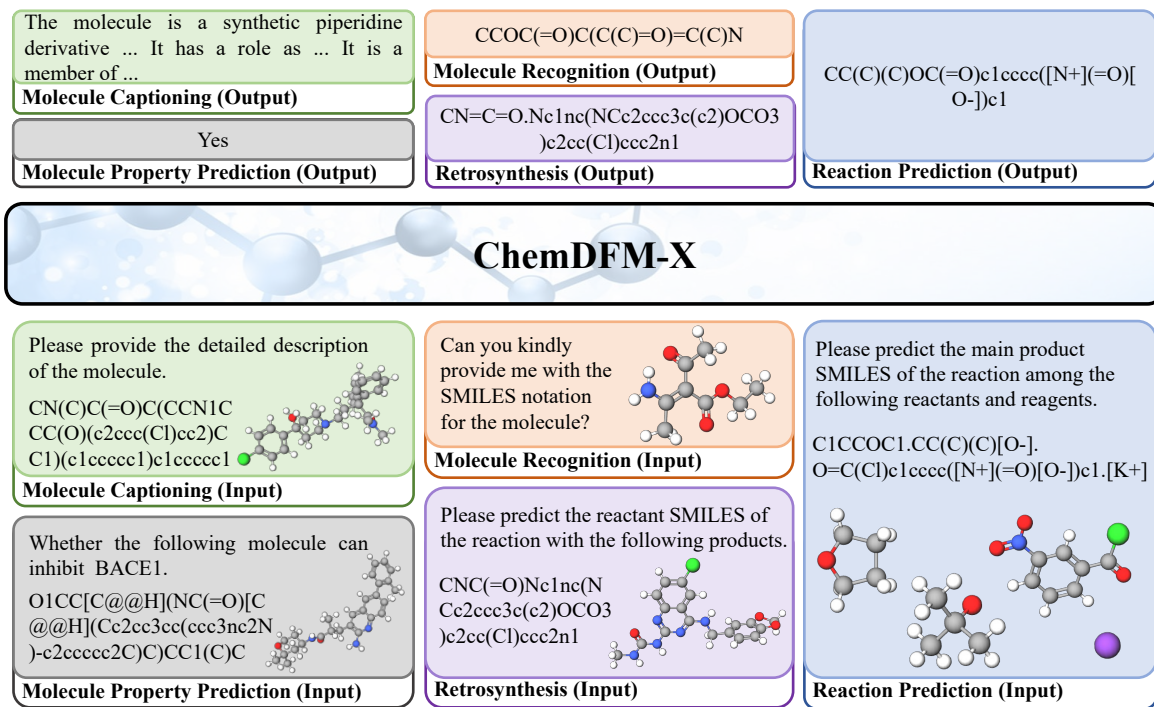


Figure B1 Evaluation tasks for structural modalities

arithmetic power. Based on the above considerations, we adopt a neural network prediction-based approach, Chemprop-IR [23], to generate IR data. We generate about 1M IR spectra in total to form the instruction-tuning dataset for training the encoder of IR spectrum modality.

IR Sequence Transformer. Similar to the situation of MS2 spectra, there is little work associated with the encoding of IR spectra. Therefore, we choose to train a new spectrum encoder for IR. We used the same setup as the mass spectrum including the model structure and the pertaining task. The data in the IR spectrum are generally denser than MS2 spectra, as the MS2 spectra data are composed of only the peaks while the IR spectra are composed of all the data points. Therefore, instead of the codebook tokenization strategy used for MS2 spectra, we directly construct the input features by reshaping the sequence of absorption intensity. After the model is trained, we use its encoder module as the encoder for the IR spectrum modality.

Appendix B Evaluation

Our evaluations are mainly conducted based on ChemLLMBench [12], one of the most widely used benchmarks designed for the evaluation of LLMs in the field of chemistry. However, considering the intrinsic differences across the modalities, the types of tasks suitable for each modality may differ. For example, the information contained in characterization modalities is often more obscure and scattered compared to that in structural modalities, therefore, it makes less sense to ask models to generate molecular captions based on the molecular spectra. Therefore, to thoroughly and specifically evaluate ChemDFM-X’s capabilities across different modalities, we construct separate evaluation task sets for each modality by both selecting existing tasks from ChemLLMBench and building new tasks. In this section, we will introduce the evaluation tasks and analyze model performances for each kind of modality separately.

Appendix B.1 Structural Modalities

Evaluation Tasks. Considering that SMILES can be viewed as a special structural modality and ChemLLMBench is originally designed for evaluating LLMs’ SMILES comprehension capabilities, all the tasks in ChemLLMBench should be suitable for the evaluation of molecular graph and molecular conformation modalities. Specifically, the evaluation tasks we utilize for both structural modalities are as follows:

- **Molecule Recognition.**⁵⁾ There are four types of molecule recognition tasks in ChemLLMBench. However, to focus tightly on evaluating non-text modality comprehension capabilities, we only source data of IUPAC to SMILES task and construct a new task where models need to recognize the molecular graphs or conformations of the molecule and give its SMILES notation.

5) Corresponding to the name prediction tasks in [12]

Model	MR*	Molecule Captioning			RP		Retrosynthesis	
	Acc	BLEU-2	BLEU-4	METEOR	Acc	Validity	Acc	Validity
<i>specialist models</i>								
Text+Chem T5 (S) [7]	-	62.5	54.2	64.8	-	-	-	-
MolCA (S + G) [20]	-	63.9	55.5	66.9	-	-	-	-
Chemformer (S) [2]	-	-	-	-	93.8	100	53.6	100
<i>generalist models</i>								
ChemLLM-7B-Chat (S) [46]	-	14.5	7.4	22.2	18.0	94.0	10.0	91.0
ChemDFM-v1.0 (S) [49]	-	26.1	20.6	33.1	49.0	98.0	12.0	99.0
3D-MoLM (S + C) [18]	0	18.2	9.9	25.6	- [†]	-	-	-
ChemDFM-X (G)	41.0	19.3	13.2	26.2	29.0	94.0	9.0	94.0
ChemDFM-X (S + G)	-	<u>26.6</u>	<u>21.1</u>	<u>33.8</u>	48.0	99.0	18.0	98.0
ChemDFM-X (C)	66.0	26.4	20.7	33.6	<u>53.0</u>	99.0	8.0	97.0
ChemDFM-X (S + C)	-	27.8	22.0	34.7	54.0	99.0	<u>16.0</u>	99.0

Table B1 The Results of multiple evaluation tasks, including Molecule Recognition (MR), Molecule Captioning, Reaction Prediction (RP), and Retrosynthesis, for the structural modalities. The metric validity evaluates whether the output SMILES is valid. The content within the parentheses indicates the molecular representation modalities used for the corresponding model (S - SMILES, G - Molecular Graphs, C - Molecular Conformations). *: The SMILES notations are not included in the inputs of this task to prevent cheating. †: 3D-MoLM does not support the reaction-related tasks as the input of 3D-MoLM can not contain more than one molecular conformation. Among the generalist models, the optimal results are **bolded**, and the second-best results are underlined.

Model	BACE	BBBP	Clintox	HIV	Tox21	Avg
<i>specialist models</i>						
Mole-BERT (G) [38]	80.8	71.9	78.9	78.2	76.8	77.3
Uni-Mol (C) [50]	85.7	72.9	91.9	80.8	79.6	82.2
MolXPT (S) [22]	88.4	80.0	95.3	78.1	77.1	83.8
MolCA (S + G) [21]	79.8	70.0	89.5	64.5	77.2	76.2
<i>generalist models</i>						
ChemLLM-7B-Chat (S) [46]	83.2	62.8	68.5	73.8	70.1	71.7
ChemDFM-v1.0 (S) [49]	78.4	66.7	<u>89.9</u>	73.6	79.8	<u>77.7</u>
3D-MoLM (S + C) [18]	50.9	48.0	47.0	53.2	49.4	49.7
ChemDFM-X (G)	66.6	65.2	45.5	70.8	79.0	65.4
ChemDFM-X (S + G)	76.1	<u>65.6</u>	92.5	<u>75.3</u>	79.3	77.8
ChemDFM-X (C)	<u>80.5</u>	65.4	71.9	75.2	79.1	74.4
ChemDFM-X (S + C)	79.0	63.3	89.5	76.4	<u>79.5</u>	77.5

Table B2 The Results of molecule property prediction tasks for the structural modalities. We report the performance of the models in AUC-ROC, which stands for the Area Under the Curve of the Receiver Operating Characteristic. The content within the parentheses indicates the molecular representation modalities used for the corresponding model (S - SMILES, G - Molecular Graphs, C - Molecular Conformations). Among the generalist models, the optimal results are **bolded**, and the second-best results are underlined.

- **Molecule Captioning.** This task requires models to generate a brief description of the given molecules. In ChemLLMBench, the molecule captioning dataset is composed of 100 instances sampled from the test set of CHEBI-20 [9]. Instead of directly using this dataset provided by ChemLLMBench, we evaluate the models on the full test set of CHEBI-20 for a more robust assessment.

- **Molecule Property Prediction.** These tasks ask models to predict the properties of the given molecules. To construct such tasks, ChemLLMBench extracts five typical classification task sets (*i.e.*, BACE, BBBP, Clintox, HIV, and Tox21) from Molecular Net [37], the widely used property prediction benchmark, and randomly sampled 100 molecule-property pairs from each task. In our evaluation, we use the same five task sets. However, to make the tasks more difficult, instead of randomly sampling, we incorporate a more challenging dataset split using the DeepChem library [27].

- **Reaction Prediction and Retrosynthesis.** These two tasks are similar in form but differ significantly in difficulty. The reaction prediction task asks models to predict the product of a chemical reaction given its reactants and reagents, which is relatively straightforward and easy. On the other hand, the retrosynthesis task asks models to predict the reactants of a reaction given its product. In our evaluation, we directly utilize the corresponding dataset provided by ChemLLMBench to evaluate these capabilities.

An intuitive diagram demonstrating the input and output of the above tasks is shown in Figure B1.

Baselines. We leverage three types of models as our baselines.

The most relevant baseline in the structure modalities is 3D-MoLM [18], which is a generalist modal that can process both SMILES and molecular conformations⁶). The primary distinctions between ChemDFM-X and 3D-MoLM can be summarized as follows: First, ChemDFM-X employs a non-compressive MLP projector, whereas 3D-MoLM uses a Q-Former module which usually leads to information losses and requires more training [41]. Second, 3D-MoLM supports only a single molecular conformation per input, preventing it from addressing tasks involving multiple molecules, such as reaction-related problems and multi-turn dialogues. In contrast, ChemDFM-X does not suffer from this intrinsic drawback as it is designed to handle the input with multiple molecular conformations in a single prompt. Furthermore, it is also worth noticing that ChemDFM-X also supports other modalities beyond molecular conformation, which 3D-MoLM does not.

The second type of baselines is other generalist models in the field of chemistry. Specifically, we use ChemDFM [49] and ChemLLM [46], two pioneer works in the field of LLM for chemistry. They excel at SMILES comprehension, but can not understand any other chemical modalities.

Finally, we also list the advanced performances of task-specific specialist models. They usually can achieve excellent performance on the targeted tasks while having zero performance on the others.

Results and Analysis. The experimental results of the structural modalities are illustrated in Table B1 and B2. The results in Table B1 show that our ChemDFM-X possesses the capability to comprehend and infer over molecular graph modality and molecular conformation modality. ChemDFM-X outperforms the baseline generalist LMM, 3D-MoLM, in all the settings while achieving better or comparable performance compared with the state-of-the-art (SOTA) performances of generalist models. Specifically, ChemDFM-X achieves the new SOTA method among the generalist models when providing both SMILES notations and molecular conformations.

However, it is worth noticing that ChemDFM-X performs relatively better with molecular conformations compared with molecular graphs. On the one hand, from a chemical perspective, molecular conformations intrinsically contain more information than molecular graphs. Therefore, the potential of molecular conformation modality as a molecular representation is greater. On the other hand, ChemDFM-X (G) only achieves 14% accuracy in the molecular recognition tasks, showing that the molecular graph modality may still remain undertrained in ChemDFM-X.

Another promising result is that the performances of ChemDFM-X in reaction-related tasks have improved significantly compared with those with only SMILES inputs. This result is consistent with the chemical intuition that molecular graphs and conformations are suitable and commonly used for chemical reaction inference.

As for the molecular property prediction tasks (Table B2), the molecular graph modality and the molecular conformation modality do not have a remarkable influence on the performances. It is worth noticing that all the involved tasks are about high-level biological properties, such as blood-brain barrier penetration and toxicity, and there exists a remarkable overlap in the information between structural modalities including SMILES. Therefore, we attribute the reason to that the molecular graphs and conformations fail to provide sufficient complementary information to enhance the molecule modeling for these tasks. Thus, to further increase the performance of these tasks, more comprehensive training or more powerful molecular representations are needed.

Appendix B.2 Image Modality

Evaluation Tasks. Considering that image modality is the most convenient modality for human researchers, our evaluations of image modality focus more on whether models can recognize the given image. Specifically, we evaluate both the molecule image recognition capability and the reaction image recognition capability of the models. The molecule image recognition task is intuitively constructed from the molecule recognition dataset used by structural modalities. As for the reaction image recognition task, we randomly sample 100 reactions from the USPTO-50K dataset [30] to form its test set. Besides, we also evaluate the molecule captioning and molecule property prediction tasks with the image modality inputs. An intuitive diagram demonstrating the input and output of the evaluation tasks for image modality is shown in Figure B2.

⁶) They also trained a 2D-MoLM which can process SMILES and molecular graphs, however, the parameter of 2D-MoLM is not yet open-sourced.

The molecule is a synthetic piperidine derivative ... It has a role as ... It is a member of ...

Molecule Captioning (Output)

Yes


Molecule Property Prediction (Output)

CCOC(=O)C(C(C)=O)=C(C)N

Molecule Image Recognition (Output)

C=C(Br)C(F)(F)F.OB(O)c1cc(Cl)cc(Cl)c1>>C=C(c1cc(Cl)cc(Cl)c1)C(F)(F)F

Reaction Image Recognition (Output)



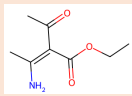
ChemDFM-X

Please provide the detailed description of the molecule.

CN(C)C(=O)C(CCN1C CC(O)(c2ccc(Cl)cc2)C C1)(c1ccccc1)c1ccccc1

Molecule Captioning (Input)

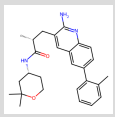
Can you kindly provide me with the SMILES notation for the molecule?



Molecule Image Recognition (Input)

Whether the following molecule can inhibit BACE1.

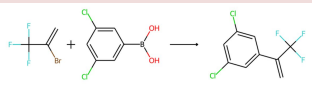
O1CC[C@@H](NC(=O)[C@@H](Cc2cc3cc(ccc3nc2N)-c2ccccc2C)CC1(C)C



Molecule Property Prediction (Input)

Please predict the reactant SMILES of the reaction with the following products.

C=C(Br)C(F)(F)F.OB(O)c1cc(Cl)cc(Cl)c1>>C=C(c1cc(Cl)cc(Cl)c1)C(F)(F)F



Reaction Image Recognition (Input)

Figure B2 Evaluation tasks for image modality

Model	Image Recognition*		Molecule Captioning		
	Molecules (Acc)	Reaction (Acc)	BLEU-2	BLEU-4	METEOR
<i>specialist models</i>					
MolNextr (I) [5]	82.0	-	-	-	-
MolScribe (I) [25]	84.0	-	-	-	-
<i>generalist models</i>					
ChemLLM-7B-Chat (S) [46]	-	-	14.5	7.4	22.2
ChemDFM-v1.0 (S) [49]	-	-	<u>26.1</u>	<u>20.6</u>	<u>33.1</u>
GPT-4O (S + I) [40]	0	0	7.0	0.4	15.0
Gemini-1.5-pro (S + I) [32]	23.0	0	8.7	2.3	14.5
Qwen-VL-max (S + I) [3]	0	0	5.9	0.1	15.0
LLaVA-v1.6-Vicuna-13B (S + I) [19]	0	0	4.4	0.0	10.7
InternLM-XComposer2 (S + I) [47]	0	0	9.2	0.0	18.3
DocOwl-1.5-Omni (S + I) [15]	0	0	10.0	1.6	17.9
ChemDFM-X (I)	46.0	53.0	23.6	18.0	30.7
ChemDFM-X (S + I)	-	-	27.2	21.7	34.3

Table B3 The Results of image recognition and molecule captioning tasks for the structural modalities. The content within the parentheses indicates the molecular representation modalities used for the corresponding model (S - SMILES, G - Molecular Graphs, I - Molecular Image). *: The SMILES notations are not included in the inputs of these tasks to prevent cheating. Among the generalist models, the optimal results are **bolded**, and the second-best results are underlined.

Model	BACE	BBBP	Clintox	HIV	Tox21	Avg
ChemLLM-7B-Chat (S) [46]	83.2	62.8	68.5	73.8	70.1	71.7
ChemDFM-v1.0 (S) [49]	<u>78.4</u>	66.7	<u>89.9</u>	73.6	79.8	<u>77.7</u>
LLaVA-v1.6-Vicuna-13B (S + I) [19]	50.0	45.5	60.6	64.6	55.0	55.1
InternLM-XComposer2 (S + I) [47]	46.8	51.7	50.3	49.6	41.7	48.0
DocOwl-1.5-Omni (S + I) [15]	44.4	53.4	50.0	57.0	48.4	50.6
ChemDFM-X (I)	69.4	<u>65.7</u>	73.3	<u>74.1</u>	77.3	72.0
ChemDFM-X (S + I)	76.9	65.1	93.5	75.3	<u>78.4</u>	77.8

Table B4 The Results of molecule property prediction tasks for the image modality. We report the performance of the models in AUC-ROC, which stands for the Area Under the Curve of the Receiver Operating Characteristic. The content within the parentheses indicates the molecular representation modalities used for the corresponding model (S - SMILES, G - Molecular Graphs, I - Molecular Image).

Baselines. Similar to structural modalities, we incorporate three types of baselines. However, there is little previous work specifically targeted at chemical images. Therefore, we leverage multiple general-domain large vision-language models as the first type of baselines, including GPT-4V [40], Gemini [32], Qwen-VL [3], LLaVA-v1.6 [19], InternLM-xcomposer2 [47], and DocOwl-1.5-Omni [15]. The second type of baselines is the same chemical LLMs as those in structural modalities, namely ChemDFM and ChemLLM. The third type of baseline is the task-specific specialist model, MolNextr [5] and MolScribe [25].

Results and Analysis. The experimental results of the image modality are illustrated in Table B3 and B4. The results show that with the help of molecule images, ChemDFM-X can achieve comparable or better performance across all the tasks, showing that our ChemDFM-X possesses strong chemical image understanding prowess.

It is worth noticing that all the general-domain LLMs we have tested fail to complete the chemical tasks. To study the reason for this phenomenon, we manually examine all the outputs of the models. We find that since these general domain LLMs are not fine-tuned for specialized chemical tasks, the models lack sufficient chemistry knowledge to complete the tasks. For instance, in the molecule captioning task, some models misinterpret it as a SMILES generation task, yielding only the SMILES notation. Others, despite comprehending the task, lack the capacity to correlate molecules with their corresponding chemical knowledge, resulting in stubbornly generating molecule descriptions unrelated to the input molecules. In the molecule image recognition task, although the models can recognize a portion of the chemical information in the image, their deficiency in understanding SMILES rules hinders them from generating long SMILES notations, limiting task performance. In the reaction image recognition task, a portion of the models claim that they don’t have the ability and refuse to answer, and other models generate invalid SMILES due to their lack of capacity for understanding the SMILES rules.

Another noteworthy result is the excellent performance of ChemDFM-X in the reaction image recognition task. Considering the length of a chemical reaction which is composed of numerous molecules, it is astonishing that the accuracy of reaction image recognition is even higher than that of single-molecule image recognition. We believe this precisely demonstrates the advantages and capabilities of cross-modality chemical LLMs, where ChemDFM-X uses the reaction knowledge it has acquired to help the reaction image recognition task. Specifically speaking, instead of single-molecule image recognition where the model needs to identify each and every bit of the molecule including the types of each atom and the connections between them, during reaction image recognition, the small mistakes when identifying each molecule may be tolerant because of the existence of the reaction context. These small mistakes may be corrected by the ChemDFM-X itself through analyzing the reaction. In other words, with the cross-modality comprehension capability, ChemDFM-X is able to utilize the reaction knowledge learned with SMILES representations to help implicitly correct minor mistakes during the reaction image recognition task, therefore resulting in much higher performances.

Appendix B.3 Characterization Modalities

Evaluation Tasks. Different from the aforementioned modalities which can only represent the known molecules, the characterization modalities are good at representing the unknown substances. The information they contained is mainly about partial properties and sub-structures and is usually implicit and scattered. Therefore, the data of characterization modalities are more suited for the unknown substance identification task.

Therefore, in our evaluation, we primarily focus on the spectrum identification tasks. Intuitively, the first group of tasks to evaluate models’ spectrum identification capabilities is basic spectrum recognition, where the spectrum of a molecule is provided and the model is asked to identify the corresponding molecule based on that. However, during real chemical research, the unknown substances do not appear out of nowhere. Usually, the spectra are provided in the context of reactions, and the knowledge such as how the substance is synthesized and what reaction it can undergo is known. Therefore, it is vital and valuable for models to possess the capability to identify the molecule based on the molecular spectra along with the reaction related to the unknown molecule. To construct such tasks, we leverage the reaction prediction and retrosynthesis tasks used for structural modalities and ask models to generate the corresponding products or reactants based on the SMILES notation of the reaction and the spectra of the missing molecule.

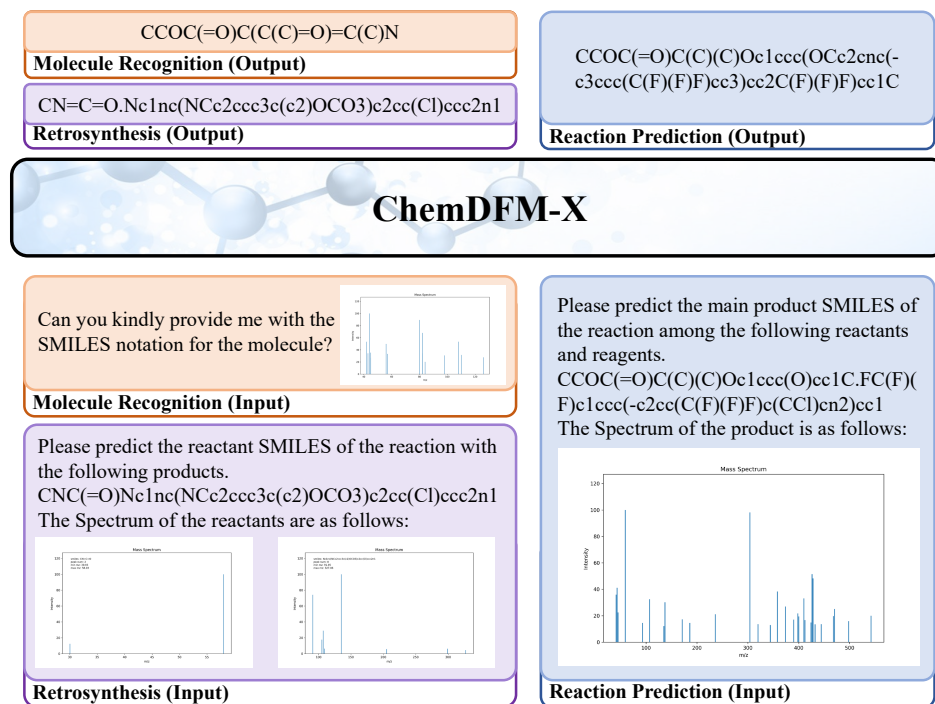


Figure B3 Evaluation tasks for characterization modalities

An intuitive diagram demonstrating the input and output of the evaluation tasks for characterization modalities is shown in Figure B3.

Baselines. The development of spectrum identification is relatively early-stage. Currently, there have still not been public-available standard baselines and benchmarks. Therefore, we only compare our ChemDFM-X model with the conventional chemical LLMs, ChemDFM, and ChemLLM, which only comprehend text and SMILES.

Results and Analysis. The experimental results of the characterization modalities are illustrated in Table B5. It is noteworthy that although the top 1 accuracy of the spectrum recognition task has near-zero performance, the performance leaps significantly when additional reaction information is involved. Note that the resulting performances (64% in the reaction prediction task and 60% in the retrosynthesis task) are far higher than those of the origin tasks where only reaction SMILES are provided with no spectra. Therefore, what is happening here is not a strong modality helping a weak modality, but a collaboration between modalities. From a chemical perspective, the spectrum modalities can provide SMILES modality hints related to the structure and composition information of the missing substances in the reactions, while SMILES modality can trim incorrect options. Through this cross-modality collaboration, ChemDFM-X is able to achieve much better performances greater than the sum of its parts. In short, ChemDFM-X successfully gains the ability to understand and analyze different spectra and is able to combine the information from different modalities to achieve higher performance.

Appendix C Related Work

AI for Chemistry has long been an active research area and has recently garnered increasing attention [4, 6, 11, 14, 16, 28, 29, 31, 38, 43–45, 48, 50]. Due to the foundational role of molecules in the chemical world, the AI for Chemistry research primarily revolves around molecules [6, 14, 16, 28, 38, 43–45, 48, 50] and reactions [4, 29, 31] which are composed of molecules. To better embed molecules into AI models, researchers mainly exploit three different molecular representations: 1) SMILES notations [6, 14, 28], 2) 2D molecular graph [38, 43, 44, 48], and 3) 3D molecular conformations [16, 45, 50]. In recent years, many works have incorporated different molecular representations together for better performance. DMP [51] encodes molecular SMILES and molecular graphs separately using a transformer and graph neural network. Through masked atom modeling and contrastive learning in both modalities, DMP pretrains the model to support inputs from both SMILES and molecular graphs. SGGRL [36] use sequence-based, graph-based, and geometry-based encoders to obtain molecular representations containing diverse modal information. Then a readout layer captures crucial information from various modalities to acquire molecule-level representations, and a fusion layer integrates different modal representations using attention mechanisms to derive the ultimate molecular representation. Although the above task-specific specialist models are able to achieve advanced performance on their corresponding tasks, they suffer from poor task generalization capabilities, and can not be applied to other tasks. Therefore, their practical utility and value is limited.

Model	SR	RP		Retrosynthesis	
	Acc	Acc	Validity	Acc	Validity
ChemLLM-7B-Chat (S) [46]	-	18	94	10	91
ChemDFM-v1.0 (S) [49]	-	<u>49</u>	<u>98</u>	12	<u>99</u>
ChemDFM-X (M)	0	-	-	-	-
ChemDFM-X (R)	1	-	-	-	-
ChemDFM-X (S + M)	-	45	88	<u>29</u>	84
ChemDFM-X (S + R)	-	64	100	60	100

Table B5 The Results of multiple evaluation tasks, including Spectrum Recognition (SR), Reaction Prediction (RP), and Retrosynthesis, for the characterization modalities. The metric validity evaluates whether the output SMILES is valid. The content within the parentheses indicates the molecular representation modalities used for the corresponding model (S - SMILES, M - MS2 Spectra, R - IR Spectra). The optimal results are **bolded**, and the second-best results are underlined.

Recently, the emerging Large Language Models (LLMs) [24,33] have shown great potential for general intelligence in the general domain. Considering the great task generalization capabilities and user interactivity of LLMs, it is also expected that LLMs can give important assistance to the research of chemistry. Both [49] and [46] explore the methods to incorporate LLMs into the field of chemistry. Specifically, [49] constructs a 34B-token domain-pretraining dataset and an instruction-tuning dataset containing 2.7M instructions. Based on these data, they specialize LLaMa-13B model [33] and get ChemDFM, one of the first chemical LLMs. ChemDFM achieves promising performances across a wide range of chemical tasks and demonstrates strong free-form dialogue capability. In addition, [46] proposes a two-stage instruction-tuning pipeline and constructs a chemical LLM called ChemLLM. These chemical LLMs significantly broaden the application scope of chemical models. However, current chemical LLMs can only process pure text and SMILES notations which are also in the form of text. Considering that chemistry is intrinsically a multi-modal subject, the capabilities achievable by text-only LLMs are limited. Most recently, [18] proposes 3D-MoLM which is the first generalist LMMs in the field of chemistry. However, their works are still in the early stages with relatively low performances and limited modalities.

In addition, Large Multimodal Models (LMMs) in the general domain have been rapidly developing. The advanced LMMs [3, 19, 32, 40, 42, 47] have achieved remarkable performance on a wide series of vision-language tasks. The promising capabilities of LMMs also enable their applications in more professional vertical domains. Taking the medical domain as an example, [34] propose Med-PaLM M, the first LMMs specialized in the medical domain, and show promising performance both in quantitative evaluation and qualitative evaluation. Additionally, [39] trained one model for each kind of modalities respectively in the medical domain. However, as the models for each modality are separate, they cannot facilitate communication and collaboration between modalities.

In this work, we construct a cross-modal chemical LLM called ChemDFM-X that can comprehend and infer chemical data of multiple different modalities with the same sets of parameters while handling a wide variety of tasks.

References

- et al H W. Chempix: automated recognition of hand-drawn hydrocarbon structures using deep learning. *Chemical Science*, 2021
- et al R I. Chemformer: a pre-trained transformer for computational chemistry. *Machine Learning: Science and Technology*, 2022
- Bai J, Bai S, Yang S, et al.. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. ArXiv:2308.12966
- Chen K, Li J, Wang K, et al. Towards an automatic ai agent for reaction condition recommendation in chemical synthesis. arXiv preprint arXiv:231110776, 2023
- Chen Y, Leung C T, Huang Y, et al. Molnextr: A generalized deep learning model for molecular image recognition. arXiv preprint arXiv:240303691, 2024
- Chithrananda S, Grand G, Ramsundar B. Chemberta: Large-scale self-supervised pretraining for molecular property prediction, 2020. ArXiv:2010.09885
- Christofidellis D, Giannone G, Born J, et al.. Unifying molecular and textual representations via multi-task language modelling, 2023. ArXiv:2301.12586
- Demartini D R. A short overview of the components in mass spectrometry instrumentation for proteomics analyses. In: Coelho A V, de Matos Ferraz Franco C, editors. *Tandem Mass Spectrometry*. : IntechOpen, 2013:
- Edwards C, Zhai C, Ji H. Text2Mol: Cross-modal molecule retrieval with natural language queries. In: Moens M F, Huang X, Specia L, et al., editors. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association: Computational Linguistics, 2021. 595-607
- EPAM. indigo-1.19.0, 2024
- Fang Y, Liang X, Zhang N, et al. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. arXiv preprint arXiv:230608018, 2023
- Guo T, Guo K, Nan B, et al. What can large language models do in chemistry? a comprehensive benchmark on eight tasks. In: *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.
- Halgren T A. Merck molecular force field. i. basis, form, scope, parameterization, and performance of mmff94. *Journal of computational chemistry*, 1996, 17: 490-519
- Honda S, Shi S, Ueda H R. Smiles transformer: Pre-trained molecular fingerprint for low data drug discovery, 2019. ArXiv:1911.04738
- Hu A, Xu H, Ye J, et al. mplug-docowl 1.5: Unified structure learning for ocr-free document understanding. arXiv preprint arXiv:240312895, 2024

- 16 Jiao R, Han J, Huang W, et al. Energy-motivated equivariant pretraining for 3d molecular graphs. In: Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence. AAAI'23/IAAI'23/EAAI'23. AAAI: Press, 2023.
- 17 Landrum G. rdkit/rdkit: 2024.03.1 (q1 2024) release. : Zenodo, 2024
- 18 Li S, Liu Z, Luo Y, et al. 3d-molm: Towards 3d molecule-text interpretation in language models. In: ICLR, 2024.
- 19 Liu H, Li C, Wu Q, et al.. Visual instruction tuning, 2023. ArXiv:2304.08485
- 20 Liu Z, Li S, Luo Y, et al. Molca: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In: EMNLP, 2023.
- 21 Liu Z, Li S, Luo Y, et al. MolCA: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In: Bouamor H, Pino J, Bali K, editors. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. Association: Computational Linguistics, 2023. 15623-15638
- 22 Liu Z, Zhang W, Xia Y, et al. MolXPT: Wrapping molecules with text for generative pre-training. In: Rogers A, Boyd-Graber J, Okazaki N, editors. Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). Association: Computational Linguistics, 2023. 1606-1616
- 23 McGill C, Forsuelo M, Guan Y, et al. Predicting infrared spectra with message passing neural networks. Journal of Chemical Information and Modeling, 2021, 61: 2594-2609
- 24 OpenAI, :, Achiam J, et al. Gpt-4 technical report. arXiv preprint arXiv:230308774, 2023
- 25 Qian Y, Guo J, Tu Z, et al. Molscribe: robust molecular structure recognition with image-to-graph generation. Journal of Chemical Information and Modeling, 2023, 63: 1925-1934
- 26 Radford A, Kim J W, Hallacy C, et al.. Learning transferable visual models from natural language supervision, 2021. ArXiv:2103.00020
- 27 Ramsundar B, Eastman P, Walters P, et al. Deep Learning for the Life Sciences. <https://www.amazon.com/Deep-Learning-Life-Sciences-Microscopy/dp/1492039837>. O'Reilly: Media, 2019
- 28 Ross J, Belgodere B, Chenthamarakshan V, et al. Large-scale chemical language representations capture molecular structure and properties. Nature Machine Intelligence, 2022, 4: 1256-1264
- 29 Sagawa T, Kojima R. Reaction5: a large-scale pre-trained model towards application of limited reaction data. arXiv preprint arXiv:231106708, 2023
- 30 Schneider N, Stiefl N, Landrum G A. What's what: The (nearly) definitive guide to reaction role assignment. Journal of Chemical Information and Modeling, 2016, 56: 2336-2346
- 31 Shi Y, Zhang A, Zhang E, et al. Relm: Leveraging language models for enhanced chemical reaction prediction. In: The 2023 Conference on Empirical Methods in Natural Language Processing, 2023.
- 32 Team G. Gemini: A family of highly capable multimodal models, 2024. ArXiv:2312.11805
- 33 Touvron H, Lavril T, Izacard G, et al.. Llama: Open and efficient foundation language models, 2023. ArXiv:2302.13971
- 34 Tu T, Azizi S, Driess D, et al.. Towards generalist biomedical ai, 2023. ArXiv:2307.14334
- 35 Wang F, Liigand J, Tian S, et al. Cfm-id 4.0: More accurate esi-ms/ms spectral prediction and compound identification. Analytical Chemistry, 2021, 93: 11692-11700
- 36 Wang Z, Jiang T, Wang J, et al. Multi-modal representation learning for molecular property prediction: Sequence, graph, geometry. ArXiv, 2024, abs/2401.03369
- 37 Wu Z, Ramsundar B, Feinberg E N, et al. Moleculenet: a benchmark for molecular machine learning. Chem Sci, 2018, 9: 513-530
- 38 Xia J, Zhao C, Hu B, et al. Mole-BERT: Rethinking pre-training graph neural networks for molecules. In: The Eleventh International Conference on Learning Representations, 2023.
- 39 Yang L, Xu S, Sellergren A, et al.. Advancing multimodal medical capabilities of gemini, 2024. ArXiv:2405.03162
- 40 Yang Z, Li L, Lin K, et al.. The dawn of lmms: Preliminary explorations with gpt-4v(ision), 2023. ArXiv:2309.17421
- 41 Yao L, Li L, Ren S, et al.. Deco: Decoupling token compression from semantic abstraction in multimodal large language models, 2024. ArXiv:2405.20985
- 42 Ye J, Hu A, Xu H, et al.. mplug-docowl: Modularized multimodal large language model for document understanding, 2023. ArXiv:2307.02499
- 43 Ying C, Cai T, Luo S, et al.. Do transformers really perform bad for graph representation?, 2021. ArXiv:2106.05234
- 44 You Y, Chen T, Sui Y, et al. Graph contrastive learning with augmentations. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS '20. Curran Associates: Inc., 2020.
- 45 Zaidi S, Schaarschmidt M, Martens J, et al. Pre-training via denoising for molecular property prediction. In: The Eleventh International Conference on Learning Representations, 2023.
- 46 Zhang D, Liu W, Tan Q, et al.. Chemllm: A chemical large language model, 2024. ArXiv:2402.06852
- 47 Zhang P, Dong X, Wang B, et al.. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition, 2023. ArXiv:2309.15112
- 48 Zhang Z, Liu Q, Wang H, et al. Motif-based graph self-supervised learning for molecular property prediction. In: Ranzato M, Beygelzimer A, Dauphin Y, et al., editors. Advances in Neural Information Processing Systems. vol. 34. Inc.: Curran Associates, 2021. 15870-15882
- 49 Zhao Z, Ma D, Chen L, et al.. Chemdfm: Dialogue foundation model for chemistry, 2024. ArXiv:2401.14818
- 50 Zhou G, Gao Z, Ding Q, et al. Uni-mol: A universal 3d molecular representation learning framework. In: The Eleventh International Conference on Learning Representations, 2023.
- 51 Zhu J, Xia Y, Qin T, et al.. Dual-view molecule pre-training, 2021. ArXiv:2106.10234