Yibo Li* Beijing University of Posts and Telecommunications Beijing, China yiboL@bupt.edu.cn

> Mengmei Zhang China Telecom Bestpay Beijing, China zhangmengmei@bestpay.com.cn

Abstract

Understanding molecular structure and related knowledge is crucial for scientific research. Recent studies integrate molecular graphs with their textual descriptions to enhance molecular representation learning. However, they focus on the whole molecular graph and neglect frequently occurring subgraphs, known as motifs, which are essential for determining molecular properties. Without such fine-grained knowledge, these models struggle to generalize to unseen molecules and tasks that require motif-level insights. To bridge this gap, we propose FineMolTex, a novel Fine-grained Molecular graph-Text pre-training framework to jointly learn coarse-grained molecule-level knowledge and fine-grained motif-level knowledge. Specifically, FineMolTex consists of two pre-training tasks: a contrastive alignment task for coarse-grained matching and a masked multi-modal modeling task for fine-grained matching. In particular, the latter predicts the labels of masked motifs and words, which are selected based on their importance. By leveraging insights from both modalities, FineMolTex is able to understand the fine-grained matching between motifs and words. Finally, we conduct extensive experiments across three downstream tasks, achieving up to 230% improvement in the text-based molecule editing task. Additionally, our case studies reveal that FineMolTex successfully captures finegrained knowledge, potentially offering valuable insights for drug discovery and catalyst design.

CCS Concepts

• Networks \rightarrow Network algorithms; • Computing methodologies \rightarrow Artificial intelligence.

*Part of the work was done as a visiting research student at Singapore Management University.

[†]Corresponding authors.

The code is publicly available at https://github.com/liushiliushi/FineMolTex.

KDD '25, Toronto, ON, Canada.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-1454-2/25/08 https://doi.org/10.1145/3711896.3736834 Yuan Fang[†] Singapore Management University Singapore yfang@smu.edu.sg

Chuan Shi[†] Beijing University of Posts and Telecommunicationsy Beijing, China shichuan@bupt.edu.cn

Keywords

Graph Neural Networks, Molecular Graph Pre-training

ACM Reference Format:

Yibo Li^{*}, Yuan Fang[†], Mengmei Zhang, and Chuan Shi[†]. 2025. Advancing Molecular Graph-Text Pre-training via Fine-grained Alignment. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '25), August 3–7, 2025, Toronto, ON, Canada.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3711896.3736834

KDD Availability Link:

The source code of this paper has been made publicly available at https://doi.org/10.5281/zenodo.15501037.

1 Introduction

Comprehending molecular structure and related knowledge is pivotal in scientific investigations spanning diverse fields, including chemistry, drug discovery, and materials science [9]. Recent advancements in artificial intelligence have yielded promising outcomes for molecule-based tasks such as retrosynthesis [33] and drug discovery [9]. The majority of these studies [1, 6, 13, 17, 27] concentrate solely on the molecular structure, such as SMILES strings, molecular graphs, and geometric structures. They learn molecular representations under supervised signals such as toxicity level and drug activity. However, this supervised learning requires extensive and costly labeling of pre-defined categories, limiting the application of previous methods to unseen categories and tasks.

Fortunately, compared to task-specific labeled data, textual descriptions of molecules are fairly abundant. These descriptions can be found in chemical database annotations, research papers in chemistry and biology, and drug instruction sheets [19], providing general information on molecular usage, efficacy, chemical properties, and even detailed insights into specific functional groups and chemical moieties [12]. Hence, several studies explore molecular structures along with their corresponding descriptions. MoleculeSTM [19] and MoMu [25] align the whole molecular graphs with their textual descriptions employing a contrastive learning approach, as shown in Figure 1(a). MoICA [20] further utilizes a cross-modal projector to map the graph embedding space to the input space of the language model. In this way, these studies reduce the reliance on task-specific labels.

However, these approaches primarily focus on the overall structure of the molecule level, failing to capture fine-grained knowledge

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '25, August 3-7, 2025, Toronto, ON, Canada.

Yibo Li*, Yuan Fang[†], Mengmei Zhang, and Chuan Shi[†]



Figure 1: Comparison of molecule- and motif-level learning, and illustration of downstream tasks.

of the sub-molecule level, such as functional groups. A natural tool to model sub-molecular structures is the motif [35], which refers to frequently recurring, significant subgraphs within molecular graphs. Motifs often play a key role in determining the properties of the whole molecular graph [35], and motif-level knowledge is frequently depicted in textual descriptions. As shown in Figure 1(b), a benzene ring indicating aromaticity property is reflected by the mention of "aromatic", and a carboxyl group is reflected by its name "carboxyl" in the description, revealing a fine-grained matching between motifs and texts.

Modeling the fine-grained motif-level knowledge is crucial for two reasons. First, motif-level knowledge is necessary for the generalization to unseen molecules, which are still largely composed of various motifs that have been seen before. For example, consider the zero-shot graph-text retrieval task shown in Figure 1(c), which aims to find the molecule most relevant to the given text. Even if the model has not been trained on the candidate molecules, it has seen many of the motifs within the unseen molecules such as the benzene and the ethyl group, corresponding to the words "benzene" and "carboxylic", respectively. Thus, the model can easily recognize the relevant molecule. Second, it bridges the gap for downstream tasks that require fine-grained knowledge. For example, in the molecule editing task illustrated in Figure 1(c), the model aims to modify part of the molecular structure based on textual instruction. This requires the model to understand the names or properties of the motifs like "chloride".

Despite the significance of this fine-grained knowledge, it is challenging to jointly learn both molecule- and motif-level knowledge, and also non-trivial to capture fine-grained matching without supervised signals. To overcome these issues, in this work, we propose a novel Fine-grained Molecular graph-Text framework (FineMolTex) to learn fine-grained motif-level knowledge, as well as coarse-grained molecule-level knowledge. First, to jointly capture both molecule- and motif-level knowledge, we use motif and word tokens to capture fine-grained knowledge, as well as two global tokens, one for the molecular graph and one for its corresponding text, to capture coarse-grained knowledge. To align this knowledge, we introduce two pre-training tasks: contrastive alignment based on global tokens and masked multi-modal learning based on motif or word tokens. Second, to capture fine-grained matching without supervised signals, as illustrated in Figure 1(b), we propose importance scores to identify important motifs and word tokens that

contain crucial fine-grained knowledge, and then selectively mask these tokens based on their importance. Finally, we incorporate a cross-attention transformer layer to integrate the embeddings of motifs and words. By predicting the labels of masked motifs and words based on information from each other, the learning of fine-grained alignment knowledge is enhanced. In summary, we outline our contributions as follows.

- We reveal that learning fine-grained motif-level knowledge provides key insight for bridging molecular graphs and text descriptions, further empowering the ability to generalize to unseen molecules and tasks.
- We introduce a novel framework named FineMolTex, consisting
 of two self-supervised pre-training tasks, to simultaneously learn
 coarse- and fine-grained knowledge. In particular, the masked
 multimodal learning task enhances the prediction for important
 masked tokens leveraging information from the other modality,
 promoting the learning of fine-grained alignment information.
- Experimental results across three downstream tasks underscore the effectiveness of FineMolTex, with a notable improvement of up to 238% in the text-based molecule editing task. Furthermore, case studies demonstrate that FineMolTex effectively aligns motifs and words, further facilitating applications such as drug discovery and catalyst design.

2 Related Work

We provide a brief review on molecular multi-modal learning. Prior works predominantly concentrate on modeling the chemical structures such as 1D SMILES [13], 2D molecular graphs [6, 17, 35], and 3D geometric structures [1, 27, 29]. They utilize supervised signals on a predetermined set, and thus cannot generalize to unseen categories without labeled examples. Recently, KV-PLM [34] bridges this gap by linking SMILES with biomedical texts through a unified language modeling framework. Nonetheless, 1D SMILES may omit certain structural details and fail to identify structural similarities among molecules due to its non-uniqueness. To address these limitations, MoleculeSTM [19] and MoMu [25] employ a contrastive learning approach to align the molecular graph with its corresponding text, thus performing well on unseen molecules and texts. However, these models are less effective on molecule-to-text generation tasks because language models are not yet well-versed in interpreting graphs as generative conditions. Therefore, MolCA [20] introduces a cross-modal projector to align the embedding

space of the molecular graph with the language model's input space, enabling the comprehension of 2D graphs as generative conditions. This approach has also been extended to 3D graph structures, where 3D-MoLM [16] uses a cross-modal projector to synchronize the embedding space of the 3D geometric structure with that of the language model. Beyond molecular graphs and textual descriptions, MV-Mol [21] further incorporates knowledge graphs, expanding the multi-modal framework to include additional sources. Additionally, various efforts have been devoted to tackling specific molecular tasks based on textual data, including zero-shot instruction molecular learning [36], molecular reaction prediction [24], and molecular relational modeling [8].

More related works on graph-based molecular learning, as well as more general multi-modal learning, can be found in Appendix A.

3 The Proposed Approach

We propose FineMolTex, a novel fine-grained molecular graph-text framework, learning both molecule- and motif-level knowledge. The model architecture is outlined in Figure 2. This section first introduces the key components in the architecture and then describes the two pre-training tasks.

3.1 Key Components of FineMolTex

To capture coarse- and fine-grained knowledge, we propose FineMol-Tex, consisting of five key components: 1) the tokenization component to decompose molecular graphs and texts into motif and word tokens; 2) a graph encoder to capture the structure of molecules and motifs; 3) a text encoder to extract the knowledge from texts and words, 4) a cross-attention layer to integrate information from different modalities; 5) a Transformer layer to generate embeddings for each token based on its contextual tokens from the same modality.

Tokenization. As shown in Figure 2, for fine-grained modeling, we fragment the molecular graphs and texts into motif tokens and word tokens. We employ the BRICS [5] algorithm to transform the molecular graph into a motif tree, and then generate a motif sequence following a breadth-first search order. Then we utilize the post-processing procedure [35] to consolidate the motif vocabulary. We break the textual description into word tokens using the word tokenizer of SciBERT [3]. For coarse-grained modeling, the global tokens of molecule and text, <MOL> and <CLS>, are inserted at the beginning of the motif and word sequences, respectively, resulting in the sequences m_0, m_1, \ldots, m_J and t_0, t_1, \ldots, t_D , where J and D are the lengths of the sequences.

Graph Encoder. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ represent a molecular graph with N atoms, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ is the set of atoms, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ denotes the bonds, and $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N] \in \mathbb{R}^{N \times \zeta}$ is the atom feature matrix. Here, \mathbf{x}_i is the feature vector of atom v_i , and ζ is the dimension of atom features. We utilize GraphMVP [18], a pre-trained Graph Isomorphism Network (GIN), to encode each motif token. GraphMVP employs multi-view pre-training to connect 2D topologies and 3D geometries, leveraging the GEOM dataset [2], which contains 250K molecular conformations. Denoting the GraphMVP encoder as f_{GraphMVP} , we encode each atom v KDD '25, August 3-7, 2025, Toronto, ON, Canada.

into an embedding as follows:

$$\mathbf{g}_{\mathbf{v}} = f_{\text{GraphMVP}}(\mathbf{x}_{\mathbf{v}}, \mathbf{x}_{\mathbf{u}}), \ u \in \mathcal{N}(v), \tag{1}$$

where $\mathcal{N}(v)$ denotes the set of neighboring atoms of v. Then we pool the atom embeddings into a motif-level embedding, $g_{\mathcal{G}}$, as follows:

$$\mathbf{g}_{\mathcal{G}} = \text{READOUT}(\{\mathbf{g}_{\mathbf{v}} | v \in \mathcal{G}\}), \quad \forall \quad \mathcal{G} \in \{m_1, m_2, \dots, m_J\}, (2)$$

where READOUT(\cdot) is permutation invariant, implemented as the average function in our model.

To preserve the intrinsic connectivity of the motifs in the original molecule, we generate position embeddings based on the breadth-first search order and incorporate them into the motif embeddings $g_{m_0}, g_{m_1}, \ldots, g_{m_J}$, resulting in updated embeddings $g'_{m_0}, g'_{m_1}, \ldots, g'_{m_J}$. **Text Encoder.** We use SciBERT [3], which has been pre-trained on texts from the chemical and biological domains, as our text encoder, denoted as f_{bert} . It can encode a text sequence as:

$$\mathbf{b}_{\mathbf{t}_0}, \mathbf{b}_{\mathbf{t}_1}, \dots, \mathbf{b}_{\mathbf{t}_D} = f_{\text{bert}}(t_0, t_1, \dots, t_D). \tag{3}$$

Subsequently, we add the position embeddings to the token embeddings following previous work [3], yielding $b'_{t_0}, b'_{t_1}, \dots, b'_{t_n}$.

Transformer Layer. To capture the contextual information for each token, we use "encoder-style" Transformer layers [28], which consist of a multi-head self-attention layer and a fully connected feed-forward network. This enables the tokens to gather information from other tokens in the same modality. We utilize f_{trm_T} and f_{trm_M} for the text and molecule modality, respectively, as follows.

$$\mathbf{z}_{t_0}, \mathbf{z}_{t_1}, \dots, \mathbf{z}_{t_D} = f_{trm_T}(\mathbf{b}'_{t_0}, \mathbf{b}'_{t_1}, \dots, \mathbf{b}'_{t_D}),$$
(4)

$$\mathbf{z}_{\mathbf{m}_{0}}, \mathbf{z}_{\mathbf{m}_{1}}, \dots, \mathbf{z}_{\mathbf{m}_{J}} = f_{\operatorname{trm}_{M}}(\mathbf{g}_{\mathbf{m}_{0}}', \mathbf{g}_{\mathbf{m}_{1}}', \dots, \mathbf{g}_{\mathbf{m}_{J}}').$$
 (5)

Cross-attention Layer. We integrate information from different modalities via cross-attention layers f_{crs_M} and f_{crs_T} for molecular graph and text, respectively. Consider the cross-attention layer f_{crs_M} for molecular graph: the queries are from the same modality, $Q_m = Z_m W_m^Q$, while the keys and values are from the text modality, $K_t = Z_t W_t^K$ and $V_t = Z_t W_t^V$. Here W_m^Q , W_t^K , W_t^V are learnable weights, $Z_m = [\mathbf{z}_{m_0}, \mathbf{z}_{m_1}, \dots, \mathbf{z}_{m_J}]$, and $Z_t = [\mathbf{z}_{t_0}, \mathbf{z}_{t_1}, \dots, \mathbf{z}_{t_K}]$. Subsequently, the output of scaled dot-product attention is computed as:

Attention
$$(Q_m, K_t, V_t) = \operatorname{softmax}\left(\frac{Q_m K_t^T}{\sqrt{d_k}}\right) V_t,$$
 (6)

where d_k is the dimension of queries and keys. The cross-attention layer for text is designed similarly. Hence, the encoding of each token accounts for tokens from the other modality, enabling the learning of fine-grained alignment at the motif level. The outputs of the cross-attention layer are:

$$\mathbf{h}_{t_0}, \mathbf{h}_{t_1}, \dots, \mathbf{h}_{t_D} = f_{crs_T}(\mathbf{z}_{t_0}, \mathbf{z}_{t_1}, \dots, \mathbf{z}_{t_D}),$$
 (7)

$$\mathbf{h}_{\mathbf{m}_0}, \mathbf{h}_{\mathbf{m}_1}, \dots, \mathbf{h}_{\mathbf{m}_J} = f_{\mathrm{crs}_M}(\mathbf{z}_{\mathbf{m}_0}, \mathbf{z}_{\mathbf{m}_1}, \dots, \mathbf{z}_{\mathbf{m}_J}). \tag{8}$$

3.2 Pre-training Tasks

We propose two pre-training tasks, the contrastive alignment task for coarse-grained alignment, and the masked multi-modal modeling task for fine-grained alignment.

KDD '25, August 3-7, 2025, Toronto, ON, Canada.

Yibo Li*, Yuan Fang[†], Mengmei Zhang, and Chuan Shi[†]



Figure 2: Architecture of FineMolTex. The input is a graph-text pair with both a molecular structure and a corresponding description. The components in the same color share the same weights.

Contrastive Alignment. For coarse-grained alignment at the molecule level, we align the graph-text pairs from the same molecules and contrast the pairs from different molecules, which can be achieved by optimizing the following loss:

$$\begin{split} L_{\rm con} &= -\frac{1}{2} \mathbb{E}_{m_0, t_0} \left[\log \frac{\exp(\cos(\mathbf{z_{m_0}}, \mathbf{z_{t_0}})/\tau)}{\exp(\cos(\mathbf{z_{m_0}}, \mathbf{z_{t_0}})/\tau) + \sum_{t_0'} \exp(\cos(\mathbf{z_{m_0}}, \mathbf{z_{t_0'}})/\tau)} \right. \\ &+ \log \frac{\exp(\cos(\mathbf{z_{t_0}}, \mathbf{z_{m_0}})/\tau)}{\exp(\cos(\mathbf{z_{t_0}}, \mathbf{z_{m_0}})/\tau) + \sum_{m_0'} \exp(\cos(\mathbf{z_{t_0}}, \mathbf{z_{m_0'}})/\tau)} \right], \end{split}$$
(9)

where $\mathbf{z}_{\mathbf{m}_0}, \mathbf{z}_{\mathbf{m}_0'}, \mathbf{z}_{\mathbf{t}_0}$, and $\mathbf{z}_{\mathbf{t}_0'}$ denote the output embeddings from the Transformer layer, t_0' and m_0' are the negative instances sampled from the same batch of graph-text pairs, and $\cos(\cdot, \cdot)/\tau$ is the cosine similarity scaled by the temperature hyperparameter τ . In this way, we capture the molecule-level knowledge, aligning the embedding space of molecular graphs and texts holistically.

Masked Multi-modal Modeling. For fine-grained alignment at the motif level, we selectively mask the important tokens and predict their labels. The token embeddings of the motifs and words are updated using ℓ_{trm_M} and ℓ_{trm_T} transformer layers, respectively. Subsequently, information from the two modalities is integrated via our cross-attention layer. This entire process is iterated for ℓ times.

Based on the output embeddings of fine-grained tokens from the cross-attention layer $\mathbf{h}_{t_1}, \ldots, \mathbf{h}_{t_D}$ and $\mathbf{h}_{m_0}, \mathbf{h}_{m_1}, \ldots, \mathbf{h}_{m_J}$, we utilize two classifiers ρ_m and ρ_t to predict the labels of the masked motifs and words: $\hat{y}_{m_i} = \rho_m(\mathbf{h}_{m_i})$, $\hat{y}_{t_j} = \rho_t(\mathbf{h}_{t_j})$, where \hat{y}_{m_i} is the predicted label of motif m_i , and \hat{y}_{t_j} is the predicted label of word t_j . Given the ground truth labels y_{m_i} and y_{t_j} , the model is trained by reconstructing the masked tokens as:

$$L_{\text{pre}} = \beta \sum_{i} \text{CE}(\hat{y}_{m_i}, y_{m_i}) + \alpha \sum_{j} \text{CE}(\hat{y}_{t_j}, y_{t_j}), \quad (10)$$

where α , β are hyperparameters, and CE(\cdot , \cdot) is the cross-entropy loss. The key to achieving fine-grained alignment lies in the crossattention layer, which enables the model to predict the labels of masked tokens based on tokens from the other modality. For instance, as illustrated in Figure 2, predicting the label of SO₃⁻ solely based on the unmasked motif tokens is challenging. However, by leveraging the embeddings of word tokens, particularly "propanesulfonic" which includes the SO₃⁻ group, we can gain relevant information about the masked token. Consequently, the model implicitly learns fine-grained alignment knowledge, thereby augmenting its motif-level knowledge.

Overall Loss. FinMolTex is optimized by the overall loss $L = L_{con} + L_{pre}$. Thus, FineMolTex is able to jointly learn the molecule- and motif-level knowledge.

3.3 Importance-based Masking

There is a large proportion of noisy tokens that lack fine-grained knowledge. For example, motif tokens such as "C", and word tokens like "an" and "this" fail to provide meaningful information. If they are used for capturing fine-grained alignment knowledge, they may negatively impact the training process.

To address this issue, we propose the importance score to identify and mask tokens that contain fine-grained alignment knowledge. Specifically, as shown in Figure 2, we observe that important motif tokens, such as m_5 , tend to contribute more to the global molecule token z_{m_0} , meaning they have a larger weight during aggregation. The same applies to the global text token z_{t_0} . Intuitively, the contrastive alignment task encourages the global token to focus more on tokens containing fine-grained knowledge, enabling better coarse-grained alignment. Thus we define the importance score as:

$$\omega_{t_i} = \frac{\exp(\cos(z_{t_i}, z_{t_0}))}{\sum_{j=1}^{N} \exp(\cos(z_{t_j}, z_{t_0}))},$$
(11)

$$\omega_{m_i} = \frac{\exp(\cos(z_{m_i}, z_{m_0}))}{\sum_{i=1}^{N} \exp(\cos(z_{m_i}, z_{m_0}))}.$$
(12)

where ω_{t_i} is the importance score of word token t_i , ω_{m_i} is the importance score of motif m_i . A higher importance score indicates a greater contribution to the global token.

We begin by only pre-training the contrastive alignment task for a few warm-up epochs to ensure valid importance scores. To demonstrate the empirical validity of the importance scores, we summarize the 10 most important and least important tokens in each modality in Tables 1 and Table 2. We observe that the top 10 word tokens are closely associated with specific motifs, indicating their relevance to fine-grained knowledge. In contrast, the bottom 10 word tokens have few associations. Similarly, the top 10 motif tokens often exhibit distinct chemical properties, making them highly meaningful, while the bottom 10 motif tokens are often less meaningful.

Then we jointly pre-train on the contrastive alignment task and masked multimodal modeling task. We mask 15% word tokens and 20% motif tokens based on their importance scores. The probability of masking each token is proportional to its score. As training progresses, the importance scores are dynamically updated, allowing the model to iteratively refine its focus on fine-grained alignment knowledge and improving overall training performance. After pretraining, the final importance scores are shown in Appendix D.

Table 1: Top 10 and bottom 10 word tokens by average importance score.

Тор 10	Bottom 10		
Token	Score	Token	Score
triglyceride	0.1419	little	0.002
peptide	0.1290	vegetables	0.002
polysaccharide	0.1284	fruits	0.002
oligo	0.1269	highest	0.002
oligomer	0.1173	game	0.002
anthocyan	0.1053	wild	0.002
diary	0.0998	meat	0.002
isofl	0.0972	chicken	0.002
cannab	0.0953	cheese	0.002
sulfide	0.093	thirty	0.002

4 Experiments

In this section, we conduct extensive experiments to demonstrate the effectiveness of FineMolTex. Before evaluating, we first conduct the two pre-training tasks on the PubChemSTM dataset [19], which includes 281K graph-text pairs from PubChem [12]. Each molecular graph is paired with a textual description that elaborates

Top 10		Bottom	Bottom 10		
Token	Score	Token	Score		
[NH4+]	0.1766	S	0.0661		
[OH-]	0.1588	С	0.0564		
C1=CN=CN=C1	0.1285	CCC	0.0521		
[NH3+]	0.1252	0	0.0503		
Cl	0.123	Ν	0.0472		
O=S	0.119	CCCC(C)OC	0.0407		
F	0.112	CCOC(C)O	0.0369		

Table 2: Top 10 and bottom 10 motif tokens by average importance score.

on its chemical and physical properties or highlights its high-level bioactivities. Details of the pre-training data and process can be found in Appendix C.1.1 and C.4.

CCOCC

C1CCOCC1

CCNC(C)=O

0.0361

0.0299

0.0231

0.108

0.1027

0.1016

C1=CC=CC=C1

Br

C1=CNC=C1

The goal of our experiments is to answer the following research questions (RQs).

RQ1. Can FineMolTex better generalize to unseen molecules?

RQ2. Can FineMolTex bridge the gap to tasks centered on motiflevel knowledge?

RQ3. Can FineMolTex perform better on single-modality tasks? **RQ4.** Has FineMolTex learned fine-grained knowledge?

RQ5. Are the token masking and cross-attention layers beneficial?

4.1 Generalization to Unseen Molecules (RQ1)

To answer RQ1, we conduct a **zero-shot graph-text retrieval task** to examine the generalizability of FineMolTex on unseen molecules and texts. Given a molecular graph and T candidate textual descriptions, the goal is to identify the textual description that best aligns with the molecular graph. Conversely, given a textual description and T candidate molecular graphs, identify the molecular graph that best matches the text. This task can be addressed by calculating the similarity of the molecular graphs and texts in the joint embedding space, thus allowing zero-shot inference.

Datasets and Baselines. We utilize DrugBank-Pharmacodynamics, molecule-ATC, and DrugBank-Description [19] extracted from the DrugBank database [31] for evaluation. These datasets include molecular graphs and their chemical descriptions. Details of the datasets can be found in Appendix C.1.2. We compare with five multimodal molecular models: KV-PLM [34], MolCA [20], MoMu-S [25], MoMu-K [25], MoleculeSTM [19], 3D-MoLM [16], and MV-Mol [21]. Specifically, KV-PLM uses SMILES to represent the structure of the molecule, while others use graph structures.

Results. We report the results on the first two datasets in Tables 3, 4, and 5. We make the following observations. 1) Across different values of T, FineMolTex consistently outperforms the baselines that neglect motif-level knowledge. The superior performance demonstrates that fine-grained motif-level knowledge facilitates generalization to unseen molecules, which likely contain seen motifs. 2) FineMolTex maintains strong performance in both directions (given graph, and given text). The symmetry further indicates that the

	Give	n Molecular G	raph	Given Text			
Т	4	10	20	4	10	20	
KV-PLM	$68.38 {\pm} 0.03$	47.59 ± 0.03	$36.54 {\pm} 0.03$	$67.68 {\pm} 0.03$	48.00 ± 0.02	34.66 ± 0.02	
MolCA	83.75 ± 0.54	74.25 ± 0.26	66.14 ± 0.21	81.27 ± 0.33	69.46 ± 0.17	62.13 ± 0.16	
MoMu-S	$70.51 {\pm} 0.04$	55.20 ± 0.15	43.78 ± 0.10	70.71 ± 0.22	54.70 ± 0.31	44.25 ± 0.43	
MoMu-K	69.40 ± 0.11	53.14 ± 0.26	42.32 ± 0.28	68.71±0.03	53.29 ± 0.05	43.83 ± 0.12	
3D-MoLM	81.35 ± 0.14	73.65 ± 0.13	64.79 ± 0.15	79.78 ± 0.22	62.38 ± 0.16	53.43 ± 0.11	
MV-Mol	92.24 ± 0.26	85.38 ± 0.19	79.41 ± 0.43	91.28 ± 0.13	85.32 ± 0.15	80.37 ± 0.22	
MoleculeSTM	92.14 ± 0.02	86.27 ± 0.02	81.08 ± 0.05	91.44 ± 0.02	86.76 ± 0.03	81.68 ± 0.03	
FineMolTex	$96.78{\pm}0.05$	$92.48{\pm}0.02$	$87.94{\pm}0.14$	$\textbf{96.29}{\pm}0.12$	$91.65{\pm}0.15$	$85.07{\pm}0.11$	

Table 3: Accuracy ($\% \pm \sigma$) of graph-text retrieval task on DrugBank-Pharmacodynamics.

Table 4: Accuracy (% $\pm \sigma$) of graph-text retrieval task on molecule-ATC.

	Give	n Molecular G	raph	Given Text		
Т	4	10	20	4	10	20
KV-PLM	$60.94 {\pm} 0.00$	42.35 ± 0.00	30.32 ± 0.00	60.67 ± 0.00	$40.19 {\pm} 0.00$	29.02±0.00
MolCA	67.34 ± 0.05	53.51 ± 0.12	44.10 ± 0.03	65.18 ± 0.34	51.01 ± 0.26	41.30 ± 0.51
MoMu-S	64.72 ± 0.04	48.72 ± 0.03	37.64 ± 0.02	64.98 ± 0.13	49.58 ± 0.05	39.04 ± 0.16
MoMu-K	61.79 ± 0.14	45.69 ± 0.22	34.55 ± 0.09	63.32 ± 0.15	47.55 ± 0.06	$37.68 {\pm} 0.18$
3D-MoLM	65.72 ± 0.08	50.48 ± 0.14	38.31 ± 0.06	63.10 ± 0.06	44.17 ± 0.11	34.56 ± 0.15
MV-Mol	70.29 ± 0.06	54.93 ± 0.14	45.64 ± 0.37	72.08 ± 0.15	59.34 ± 0.22	48.56 ± 0.36
MoleculeSTM	69.33 ± 0.03	54.83 ± 0.04	44.13 ± 0.05	71.81 ± 0.05	58.34 ± 0.07	$47.58 {\pm} 0.05$
FineMolTex	$76.52{\pm}0.10$	$62.75{\pm}0.06$	$51.84{\pm}0.16$	$76.38{\pm}0.18$	$61.72{\pm}0.09$	$\textbf{50.88}{\pm}0.13$

Table 5: Accuracy ($\% \pm \sigma$) of graph-text retrieval task on DrugBank-Description.

	Give	n Molecular G	raph	Given Text		
Т	4	10	20	4	10	20
KV-PLM	$73.80 {\pm} 0.00$	53.96 ± 0.29	40.07 ± 0.38	72.86 ± 0.00	52.55 ± 0.29	$40.33 {\pm} 0.00$
MolCA	93.75 ± 0.09	87.25 ± 0.06	82.77 ± 0.12	90.71 ± 0.04	84.97 ± 0.16	77.53 ± 0.15
MoMu-S	76.52 ± 0.12	61.66 ± 0.25	50.00 ± 0.08	77.62 ± 0.06	61.49 ± 0.15	52.20 ± 0.13
MoMu-K	74.15 ± 0.08	57.18 ± 0.16	47.97 ± 0.14	77.79 ± 0.12	62.33 ± 0.18	47.97 ± 0.06
3D-MoLM	92.81±0.23	85.71±0.19	80.20 ± 0.33	88.31 ± 0.32	81.23 ± 0.07	$74.40 {\pm} 0.39$
MV-Mol	95.13 ± 0.16	90.28 ± 0.21	84.83 ± 0.34	93.54 ± 0.32	86.58 ± 0.52	80.75 ± 0.43
MoleculeSTM	99.15 ± 0.00	97.19 ± 0.00	95.66 ± 0.00	99.05 ± 0.37	97.50 ± 0.46	95.71 ± 0.46
FineMolTex	$99.60{\pm}0.06$	$97.96{\pm}0.04$	$\textbf{96.70}{\pm}0.14$	$99.62{\pm}0.02$	$97.96{\pm}0.09$	$96.34{\pm}0.07$

embedding spaces of both modalities are well-aligned and similarly well-learned. 3) We observe that KV-PLM, which utilizes SMILES to capture molecular structures, is less effective than other models employing graphs, consistent with previous findings [19] that 2D graph structure is more expressive than 1D SMILES.

4.2 Application to Motif-Centered Tasks (RQ2)

To answer RQ2, we employ a zero-shot **text-based molecule editing task**, which is highly relevant to practical applications including catalyst design and targeted drug discovery. Specifically, we utilize FineMolTex to collaborate with a molecule generation module, following the design in [19], to modify a specified molecule according to a text prompt. Hence, motif-level knowledge is essential for this task, as the model needs to replace certain motifs with others that are related to specific properties and names as indicated in the text prompt. We defer the technical details to Appendix B. We randomly sample 200 molecules from ZINC [11], and select 12 text prompts, including 8 prompts pertaining to physical properties [19], and 4 based on the names of the motifs. We utilize MoleculeSTM, MoMu, and MolCA as baselines.

Evaluation. We employ different methods to assess whether the generated molecules satisfy the two types of prompts. For the 8 prompts on physical properties, we employ three measures: LogP, QED, and tPSA, which measures solubility [15], drug-likeness [4],



Figure 3: Hit ratios of 12 text-based molecule editing tasks.



Figure 4: Visual analysis of the output molecules of MoleculeSTM and Motif-MolTex on 4 text-based molecule editing tasks. Differences between the input and output molecule of FineMolTex are highlighted in red and green circles. Lower LogP indicates higher water solubility.

and permeability [7], respectively. We consider the editing to be successful if the difference in measurements between the input and output molecules exceeds a specified threshold Δ , which we have set to 0 following one of the settings in literature [19]. For the 4 prompts based on motif names, we use RDKit [14] to verify the presence of the indicated motifs. For all 12 prompts, we report the *hit ratio*: the proportion of generated molecules that meet our expectations.

Results. As shown in Figure 3, FineMolTex shows superior performance on these prompts, especially on the 4 prompts with motif names. Notably, we achieve a relative gain of up to 230% over the

best-performing baseline, demonstrating that FineMolTex has an advanced understanding of motif-level knowledge. We also visualize the output molecules of MoleculeSTM and FineMolTex in Figure 4. It can be observed that while MoleculeSTM produces incorrect molecules, FineMolTex accurately generates the intended molecules. Specifically, when prompted to generate molecules that are soluble in water, FineMolTex successfully creates molecules with lower LogP than the input molecule, as lower LogP indicates higher water solubility. Similarly, when prompted to generate molecules with hydroxyl groups or chlorine atoms, FineMolTex correctly does

Yibo Li*, Yuan Fang[†], Mengmei Zhang, and Chuan Shi[†]

so. These results confirm that FineMolTex possesses a deeper understanding of motif-level knowledge, thereby enhancing the generative capabilities.

4.3 Application to Single-Modality Task (RQ3)

While FineMolTex can simultaneously utilize pre-trained knowledge from both graphs and texts, we also verify its effectiveness on single-modality tasks, namely, **molecular property prediction tasks**. We use MoleculeNet [32] as the dataset, which only provides molecular graphs as input without texts. More specifically, there are eight binary classification tasks, and we report ROC-AUC for evaluation. More detailed dataset descriptions are provided in Appendix C.1.3.

Baselines. We compare FineMolTex against nine baselines, including 1) five pre-trained GNN models: AttrMasking [10], ContextPred [10], InfoGraph [26], MolCLR [30], and GraphMVP [18]; 2) three graph-text multimodal models: MoMu-S [25], MoMu-K [25], and MoleculeSTM [19], and 3) one SMILES-text multimodal model: KV-PLM [34].

Results. As shown in Table 6, FineMolTex consistently outperforms all baselines, achieving relative gains of 3.2%, 2.4%, and 4.7% on SIDER, MUV, and BACE, respectively, compared to the best baseline. The promising performance of FineMolTex indicates that it implicitly utilizes pre-trained knowledge from the text modality even when the input consists solely of graphs. Additionally, KV-PLM exhibits a notable performance gap from other models, due to its use of 1D SMILES strings for molecular structure and a smaller pre-training dataset.

4.4 Analysis of Learned Fine-grained Knowledge (RQ4)

We evaluate whether FineMolTex captures fine-grained alignment information in the joint embedding space, and assess if it can predict the labels of masked motifs based on fine-grained knowledge.

Visualization of Motif and Word Embeddings. To evaluate whether FineMolTex captures fine-grained alignment knowledge, we select motif and word tokens from 1,500 graph-text pairs in the PubChemSTM dataset, excluding meaningless words such as "this" and "a". In total, we visualize 3,469 motif tokens and 6,914 text tokens with *t*-SNE [22] in Figure 5, where triangles denote text tokens, and circles denote motif tokens, with different colors indicating various labels. To examine the details of the tokens, we zoom into several regions in the figure, retaining only the colors and legends of the tokens we are interested in. For brevity, we utilize SMILES to represent the motif structures. We observe that text and motif tokens corresponding to each other are also close in the embedding space. For instance, in the pink frame, the word "ammonium" is close to the motif tokens "[NH2+]=O", "C=[NH2+]", and "[NH3+]O", which are related to "ammonium." In the blue frame, the word "poison" is adjacent to the motifs "[AsH3]" and "O[AsH2]", which are poisonous. In the orange frame, the word "sulf" is close to the motif tokens "OS", "CCSSCC", and "CC1=CSC(C)=N1", all of which represent sulfides. These results demonstrate that FineMolTex learns the connections between motifs and their chemical names and properties, thereby significantly enhancing its expressiveness.



Figure 5: Visualization of motif tokens and word tokens using *t*-SNE. Triangles denote word tokens; circles denote motif tokens.



Figure 6: Explaination of the prediction of certain masked motifs based on text tokens utilizing LIME.

Predictions Based on Fine-grained knowledge. To further verify that FineMolTex can utilize the learned fine-grained alignment knowledge for predictions, we utilize Local Interpretable Model-Agnostic Explanation (LIME) [23], a well-established tool that can explain the predictions of certain masked motifs based on text tokens. By perturbing the input text, LIME observes how the model's predictions change with variations in the input text. Then, LIME fits these perturbed texts and the prediction results to an interpretable model such as a linear model. This approach allows us to quantify the significance of each text token in predicting the motifs, thereby revealing the fine-grained knowledge learned by FineMolTex. The results are shown in Figure 6, where text tokens with higher interpretive weights are more crucial for predictions, and thus more relevant to the masked motifs. Specifically, the word with the highest interpretive weights in (1) is "heteroarene", which directly refers to the name of the masked motif. In (2), the word "phenyl" refers to the masked motif. In (3), the masked motif makes the molecule "bromide". These findings demonstrate that FineMol-Tex has effectively acquired motif-level knowledge.

Model	BBBP	Tox21	ToxCast	Sider	ClinTox	MUV	HIV	Bace	Avg
AttrMask	67.8±2.6	$75.0 {\pm} 0.2$	$63.6 {\pm} 0.8$	58.1±1.2	75.4 ± 8.8	73.8±1.2	75.4 ± 0.5	80.3±0.0	71.2
ContextPred	63.1±3.5	74.3 ± 0.2	$61.6 {\pm} 0.5$	60.3 ± 0.8	80.3 ± 3.8	71.4 ± 1.4	70.7 ± 3.6	$78.8 {\pm} 0.4$	70.1
InfoGraph	64.8 ± 0.6	76.2 ± 0.4	62.7 ± 0.7	59.1±0.6	76.5 ± 7.8	73.0 ± 3.6	70.2 ± 2.4	77.6 ± 2.0	70.0
MolCLR	67.8 ± 0.5	$67.8 {\pm} 0.5$	64.6 ± 0.1	58.7 ± 0.1	84.2 ± 1.5	72.8 ± 0.7	75.9 ± 0.2	71.1 ± 1.2	71.3
GraphMVP	68.1±1.4	77.1 ± 0.4	$65.1 {\pm} 0.3$	60.6 ± 0.1	84.7 ± 3.1	$74.4 {\pm} 2.0$	77.7 ± 2.5	80.5 ± 2.7	73.5
GraphCL	69.7 ± 0.7	$73.9 {\pm} 0.7$	$62.4 {\pm} 0.6$	60.5 ± 0.9	76.0 ± 2.7	69.8 ± 2.7	78.5 ± 1.2	75.4 ± 1.4	70.8
KV-PLM	70.5 ± 0.5	72.1 ± 1.0	55.0 ± 1.7	59.8 ± 0.6	89.2 ± 2.7	54.6 ± 4.8	65.4 ± 1.7	78.5 ± 2.7	68.2
MoMu-S	70.5 ± 2.0	75.6 ± 0.3	$63.4 {\pm} 0.5$	60.5 ± 0.9	79.9 ± 4.1	70.5 ± 1.4	$75.9 {\pm} 0.8$	76.7 ± 2.1	71.6
MoMu-K	70.1 ± 1.4	$75.6 {\pm} 0.5$	$63.0 {\pm} 0.4$	60.4 ± 0.8	77.4 ± 4.1	71.1 ± 2.7	76.2 ± 0.9	77.1 ± 1.4	71.4
MolCA	$70.0 {\pm} 0.5$	$77.2{\pm}0.5$	64.5 ± 0.8	63.0 ± 1.7	89.5 ± 0.7	72.1 ± 1.3	77.2 ± 0.6	$79.8 {\pm} 0.5$	74.2
MoleculeSTM	70.0 ± 0.5	76.9 ± 0.5	$65.1 {\pm} 0.4$	61.0 ± 1.1	92.5 ± 1.1	73.4 ± 2.9	77.0 ± 1.8	80.8 ± 1.3	74.6
FineMolTex	$73.5{\pm}1.6$	77.1 ± 1.2	$68.6{\pm}0.9$	$64.8{\pm}1.4$	$92.5{\pm}0.8$	$76.3{\pm}1.2$	$79.0{\pm}1.4$	$84.0{\pm}1.5$	76.9

Table 6: Downstream results ($\% \pm \sigma$) on eight binary classification datasets from MoleculeNet.

Table 7: Ablation study ($\% \pm \sigma$) on molecule-ATC and DrugBank-Pharmacodynamics.

	molecule-AT	2	DrugBank-Pharmacodynamics		
	Given Molecular Graph Given Text		Given Molecular Graph	Given Text	
w/o mmm	68.85 ± 0.32	$69.34 {\pm} 0.14$	$90.18 {\pm} 0.08$	90.52 ± 0.14	
motif mask only	72.64 ± 0.05	$70.96 {\pm} 0.20$	92.24 ± 0.12	92.06 ± 0.31	
word mask only	73.68 ± 0.11	$72.47 {\pm} 0.09$	93.28 ± 0.26	92.97 ± 0.17	
w/o cross-attention	69.92 ± 0.25	69.35 ± 0.22	92.66 ± 0.08	92.85 ± 0.24	
random mask only	73.95 ± 0.15	$73.34 {\pm} 0.19$	93.57 ± 0.32	93.25 ± 0.18	
FineMolTex	$76.52 {\pm} 0.10$	$76.38{\pm}0.18$	96.78±0.05	$\textbf{96.29}{\pm}0.12$	

Additional analysis such as fine-grained alignment information in the joint embedding space can be found in Appendix D.3.

4.5 Ablation Study for Masking and Cross-Attention Layers (RQ5)

To thoroughly explore the impact of the key components in FineMol-Tex, we compare to several variants, including w/o mmm, which drops the masked multimodal modeling task altogether; motif mask only, which only mask motif tokens; word mask only, which only mask word tokens; w/o cross-attention, which excludes cross-attention layers. random mask only, which masks motif or word tokens randomly rather than based on importance. We evaluate these variants on the graph-text retrieval task used in RQ1, on two datasets with T = 4. As reported in Table 7, FineMol-Tex consistently surpasses the other variants. Specifically, without the masked multimodal modeling task, "w/o mmm" fails to capture fine-grained knowledge at all, resulting in the poorest performance. "motif mask only" and "word mask only" outperform "w/o mmm," because they still enable some level of fine-grained knowledge learning by either predicting motifs based on word tokens or predicting word tokens based on motif tokens. However, they are less effective than FineMolTex, which masks both words and motifs for mutual alignment. Lastly, without the cross-attention layers, "w/o cross-attention" cannot integrate token embeddings from different modalities, without importance-based masking, "random mask only" cannot effectively focus on the tokens that contain finegrained knowledge, thereby hampering their ability to effectively

learn fine-grained knowledge. These observations demonstrate the effectiveness of each component.

Additional experimental results, including further ablation studies and a comparison of pre-training and inference times, are provided in Appendix D.

5 Conclusions

In this paper, we reveal that fine-grained motif-level knowledge is crucial for molecular representation learning. We propose FineMol-Tex to jointly learn both coarse- and fine-grained knowledge through a contrastive alignment task and a masked multimodal learning task, respectively. By selectively masking the important motif/word tokens and predicting their labels using tokens from the other modality, we can effectively learn fine-grained alignment between motifs and words. Experimental results on three downstream tasks and two case studies demonstrate the effectiveness of FineMolTex.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (No. 62192784, U22B2038, 62472329). This research is also supported by the Ministry of Education, Singapore under its Academic Research Fund (AcRF) Tier 1 grant (22-SIS-SMU-054). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore. KDD '25, August 3-7, 2025, Toronto, ON, Canada.

References

- Kenneth Atz, Francesca Grisoni, and Gisbert Schneider. 2021. Geometric deep learning on molecular representations. *Nat. Mach. Intell.* 3, 12 (2021), 1023–1032. doi:10.1038/S42256-021-00418-8
- [2] Simon Axelrod and Rafael Gómez-Bombarelli. 2022. GEOM, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data* (Apr 2022). doi:10.1038/s41597-022-01288-4
- [3] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A pretrained language model for scientific text. arXiv preprint arXiv:1903.10676 (2019).
- [4] G. Richard Bickerton, Gaia V. Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L. Hopkins. 2012. Quantifying the chemical beauty of drugs. *Nature Chemistry* (Feb 2012), 90–98. doi:10.1038/nchem.1243
- [5] Jörg Degen, Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. 2008. On the art of compiling and using "drug-like" chemical fragment spaces. *ChemMedChem* (Oct 2008), 1503–1507. doi:10.1002/cmdc.200800178
- [6] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints. In Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (Eds.). 2224–232. https://proceedings.neurips.cc/paper/2015/hash/f9be311e65d81a9ad8150a60844bb94c-Abstract.html
- [7] Peter Ertl, Bernhard Rohde, and Paul Selzer. 2000. Fast Calculation of Molecular Polar Surface Area as a Sum of Fragment-Based Contributions and Its Application to the Prediction of Drug Transport Properties. *Journal of Medicinal Chemistry* 43, 20 (Oct 2000), 3714–3717. doi:10.1021/jm000942e
- [8] Junfeng Fang, Shuai Zhang, Chang Wu, Zhengyi Yang, Zhiyuan Liu, Sihang Li, Kun Wang, Wenjie Du, and Xiang Wang. 2024. MolTC: Towards Molecular Relational Modeling In Language Models. *CoRR* abs/2402.03781 (2024). doi:10. 48550/ARXIV.2402.03781 arXiv:2402.03781
- [9] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *International* conference on machine learning. PMLR, 1263–1272.
- [10] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2019. Strategies for pre-training graph neural networks. arXiv preprint arXiv:1905.12265 (2019).
- [11] John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. 2012. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling* 52, 7 (2012), 1757–1768.
- [12] Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A Shoemaker, Paul A Thiessen, Bo Yu, et al. 2021. Pub-Chem in 2021: new data content and improved web interfaces. *Nucleic acids research* 49, D1 (2021), D1388–D1395.
- [13] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alán Aspuru-Guzik. 2020. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach. Learn. Sci. Technol.* 1, 4 (2020), 45024. doi:10.1088/2632-2153/ABA947
- [14] G Landrum. 2024. Rdkit: Open-source cheminformatics. http://www.rdkit.org
 [15] Leo Leo, Albert Albert, Hansch Hansch, Corwin Corwin, Elkins Elkins, and David
- David. 1971. Partition coefficients and their uses. (Jan 1971).
- [16] Sihang Li, Zhiyuan Liu, Yanchen Luo, Xiang Wang, Xiangnan He, Kenji Kawaguchi, Tat-Seng Chua, and Qi Tian. 2024. Towards 3D Molecule-Text Interpretation in Language Models. *CoRR* abs/2401.13923 (2024). doi:10.48550/ARXIV.2401.13923 arXiv:2401.13923
- [17] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. 2019. N-Gram Graph: Simple Unsupervised Representation for Graphs, with Applications to Molecules. In Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 8464–8476. https://proceedings.neurips.cc/paper/2019/hash/2f3926f0a9613f3c3cc21d52a3cdb4d9-Abstract.html
- [18] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. 2019. N-Gram Graph: Simple Unsupervised Representation for Graphs, with Applications to Molecules. In Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 8464–8476. https://proceedings.neurips.cc/paper/2019/hash/2f3926f0a9613f3c3cc21d52a3cdb4d9-Abstract.html
- [19] Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Animashree Anandkumar. 2023. Multi-modal molecule structure-text model for text-based retrieval and editing. *Nature Machine Intelligence* 5, 12 (2023), 1447–1457.
- [20] Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. 2023. MolCA: Molecular Graph-Language Modeling

with Cross-Modal Projector and Uni-Modal Adapter. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 15623–15638. doi:10.18653/V1/2023. EMNLP-MAIN.966

- [21] Yizhen Luo, Kai Yang, Massimo Hong, Xing Yi Liu, Zikun Nie, Hao Zhou, and Zaiqing Nie. 2024. Learning Multi-view Molecular Representations with Structured and Unstructured Knowledge. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2024, Barcelona, Spain, August 25-29, 2024, Ricardo Baeza-Yates and Francesco Bonchi (Eds.). ACM, 2082–2093. doi:10.1145/3637528.3672043
- [22] Laurensvander Maaten and GeoffreyE. Hinton. 2008. Visualizing Data using t-SNE. Journal of Machine Learning Research, Journal of Machine Learning Research (Jan 2008).
- [23] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 1135–1144. doi:10.1145/2939672.2939778
- [24] Yaorui Shi, An Zhang, Enzhi Zhang, Zhiyuan Liu, and Xiang Wang. 2023. ReLM: Leveraging Language Models for Enhanced Chemical Reaction Prediction. In Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 5506–5520. doi:10.18653/V1/2023.FINDINGS-EMNLP.366
- [25] Bing Su, Dazhao Du, Zhao Yang, Yujie Zhou, Jiangmeng Li, Anyi Rao, Hao Sun, Zhiwu Lu, and Ji-Rong Wen. 2022. A Molecular Multimodal Foundation Model Associating Molecule Graphs with Natural Language. *CoRR* abs/2209.05481 (2022). doi:10.48550/ARXIV.2209.05481 arXiv:2209.05481
- [26] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. 2019. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. arXiv preprint arXiv:1908.01000 (2019).
- [27] Artur P. Toshev, Gianluca Galletti, Johannes Brandstetter, Stefan Adami, and Nikolaus A. Adams. 2023. E(3) Equivariant Graph Neural Networks for Particle-Based Fluid Mechanics. *CoRR* abs/2304.00150 (2023). doi:10.48550/ARXIV.2304. 00150 arXiv:2304.00150
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- [29] Ruijia Wang, YiWu Sun, Yujie Luo, Shaochuan Li, Cheng Yang, Xingyi Cheng, Hui Li, Chuan Shi, and Le Song. 2023. Injecting Multimodal Information into Rigid Protein Docking via Bi-level Optimization. In Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/ 77fa0e7d45c6687f1958de0b31e9fc05-Abstract-Conference.html
- [30] Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. 2021. MolCLR: Molecular Contrastive Learning of Representations via Graph Neural Networks. CoRR abs/2102.10056 (2021). arXiv:2102.10056 https://arxiv.org/abs/ 2102.10056
- [31] David S. Wishart, Yannick D. Feunang, An Chi Guo, Elvis J. Lo, Ana Marcu, Jason R. Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, Nazanin Assempour, Ithayavani Jynkkaran, Yifeng Liu, Adam Maciejewski, Nicola Gale, Alex Wilson, Lucy Chin, Ryan Cummings, Diana Le, Allison Pon, Craig Knox, and Michael Wilson. 2018. DrugBank 5.0: a major update to the DrugBank database for 2018. Nucleic Acids Res. 46, Database-Issue (2018), D1074–D1082. doi:10.1093/NAR/GKX1037
- [32] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science* 9, 2 (2018), 513–530.
- [33] Chaochao Yan, Qianggang Ding, Peilin Zhao, Shuangjia Zheng, Jinyu Yang, Yang Yu, and Junzhou Huang. 2020. Retroxpert: Decompose retrosynthesis prediction like a chemist. Advances in Neural Information Processing Systems 33 (2020), 11248–11258.
- [34] Zheni Zeng, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2022. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications* 13, 1 (2022), 862.
- [35] Zaixi Zhang, Qi Liu, Hao Wang, Chengqiang Lu, and Chee-Kong Lee. 2021. Motifbased Graph Self-Supervised Learning for Molecular Property Prediction. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 15870–15882. https://proceedings.neurips. cc/paper/2021/hash/85267d349a5e647ff0a9edcb5ffd1e02-Abstract.html
- [36] Haiteng Zhao, Shengchao Liu, Chang Ma, Hannan Xu, Jie Fu, Zhihong Deng, Lingpeng Kong, and Qi Liu. 2023. GIMLET: A Unified Graph-Text Model for

Instruction-Based Molecule Zero-Shot Learning. In Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/ 129033c7c08be683059559e8d6bfd460-Abstract-Conference.html