

3D-MoLM: TOWARDS 3D MOLECULE-TEXT INTERPRETATION IN LANGUAGE MODELS

Sihang Li^{1,3*}Zhiyuan Liu^{2*}Yanchen Luo^{1,3}Xiang Wang^{1,4†}Xiangnan He^{1,3†}Kenji Kawaguchi²Tat-Seng Chua²Qi Tian⁵¹University of Science and Technology of China ²National University of Singapore³MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition, USTC⁴Institute of Dataspace, Hefei Comprehensive National Science Center ⁵Huawei Cloud

ABSTRACT

Language Models (LMs) have greatly influenced diverse domains. However, their inherent limitation in comprehending 3D molecular structures has considerably constrained their potential in the biomolecular domain. To bridge this gap, we focus on 3D molecule-text interpretation, and propose **3D-MoLM**: 3D-Molecular Language Modeling. Specifically, 3D-MoLM enables an LM to interpret and analyze 3D molecules by equipping the LM with a 3D molecular encoder. This integration is achieved by a 3D molecule-text projector, bridging the 3D molecular encoder’s representation space and the LM’s input space. Moreover, to enhance 3D-MoLM’s ability of cross-modal molecular understanding and instruction following, we meticulously curated a 3D molecule-centric instruction tuning dataset – **3D-MoIT**. Through 3D molecule-text alignment and 3D molecule-centric instruction tuning, 3D-MoLM establishes an integration of 3D molecular encoder and LM. It significantly surpasses existing baselines on downstream tasks, including molecule-text retrieval, molecule captioning, and more challenging open-text molecular QA tasks, especially focusing on 3D-dependent properties. We release our codes and datasets at <https://github.com/lsh0520/3D-MoLM>.

1 INTRODUCTION

The advancement of Language Models (LMs) (Devlin et al., 2019; OpenAI, 2023b; Touvron et al., 2023a) has triggered a series of remarkable innovations across multiple disciplines (Zhao et al., 2023). Notably, LMs excel at text-based molecule understanding tasks, such as question-answering (QA) in the chemical and medical domains (Taylor et al., 2022), by pretraining on extensive biochemical literature. Recognizing the potential of LMs in harnessing extensive biochemical knowledge for molecule-relevant tasks, molecule-text modeling emerges as a new research direction (Edwards et al., 2021; 2022). Previous works have been dedicated to harmonizing texts with 1D molecular sequences (Zeng et al., 2022; Taylor et al., 2022) and 2D molecular graphs (Liu et al., 2023b; Su et al., 2022; Liu et al., 2022a), aiding in tasks like molecule-text retrieval and molecule captioning. However, they mostly leave 3D molecular structures untouched, which are crucial to understanding molecular dynamics, protein-ligand interactions, enzymatic functions, and a range of other biomolecular phenomena (Karplus & McCammon, 2002; Jorgensen, 2004).

To bridge this gap, we focus on 3D molecule-text interpretation, with the goal of enabling an LM to interpret and analyze 3D molecular structures through text generation. Given the recent successes of 3D molecular encoders in tasks like molecule property prediction, docking, and conformation prediction (Zhou et al., 2023; Lu et al., 2023; Fang et al., 2022), it is promising to incorporate one as an LM’s perception module for 3D molecules. Upon examination of existing literature (Dai et al., 2023; Hong et al., 2023; Chung et al., 2022), we identify two key challenges to seamlessly integrate a 3D molecular encoder into an LM for 3D molecule-text interpretation:

- **3D Molecule-Text Alignment** maps 3D molecular representations into the input textual space where the LM can understand.

*Equal contribution.

†Correspondence to Xiang Wang and Xiangnan He. {xiangwang1223, xiangnanhe}@gmail.com

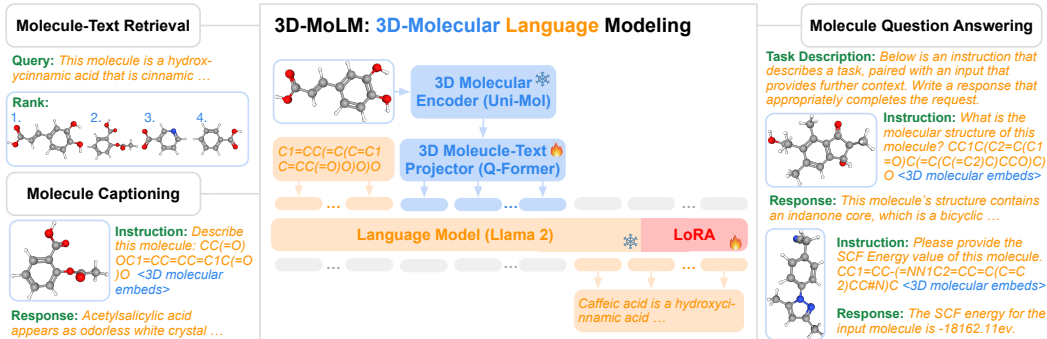


Figure 1: Demonstration of 3D-MoLM. 3D-MoLM is a general-purpose molecular LM that can be applied for molecule-text retrieval, molecule captioning, and molecular QA tasks. Flame 🔥 denotes tunable modules, while snowflake ❄️ indicates frozen modules.

- **3D Molecule-centric Instruction Tuning** fine-tunes the model to follow human instructions on 3D molecule relevant tasks.

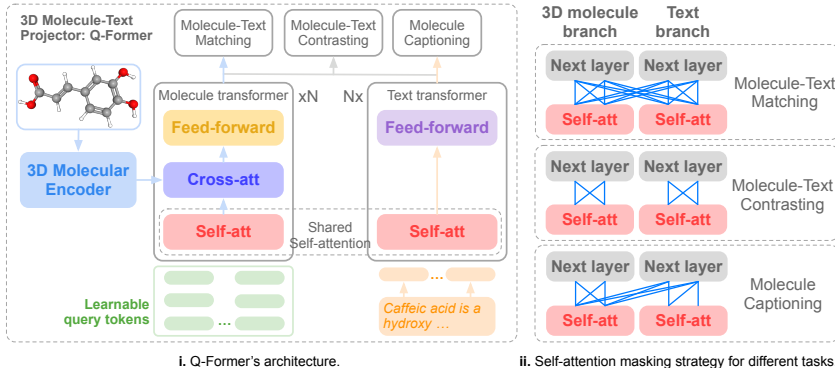
To address these challenges, we propose **3D-MoLM: 3D-Molecular Language Modeling**, as depicted in Figure 1. Specifically, it consists of two key components: 1) a 3D molecule-text projector for 3D molecule-text alignment, which aligns the latent representation spaces between the 3D molecular encoder and the LM, and 2) a dataset for **3D Molecule-centric Instruction Tuning**, **3D-MoIT**, as shown in Figure 3. 3D-MoIT enhances the model’s ability to follow human instructions and discern 3D-dependent properties of molecules.

For 3D molecule-text alignment, we employ Q-Former (Li et al., 2023) as the 3D molecule-text projector, drawing inspiration from leading vision-language modeling methods (Zhu et al., 2023; Dai et al., 2023). Given a molecule’s 3D structure, Q-Former converts it into tokens, which serve as 1D soft prompts (Li & Liang, 2021), harmonizing seamlessly with the language space of the LM. This translation facilitates the LM’s interpretation of 3D molecular structures. To cultivate the Q-Former’s alignment capability, two training stages are conducted – the first stage focuses on 3D molecule-text representation learning, while the second stage optimizes for 3D molecule-text alignment. As depicted in Figure 3, these two training stages are facilitated by our collected 316K molecule-text pairs from PubChem (Kim et al., 2021). To promote the 3D molecule-text alignment process, we manipulate the dataset by generating the 3D conformations based on SMILES using RDKit (Landrum et al., 2013) and enriching the molecular descriptions with GPT-3.5 (OpenAI, 2023a). We will detail the collection and enrichment of PubChem Dataset in Section 2.2.1 and Appendix B.


Upon aligning 3D molecules with texts, we conduct instruction tuning using our curated dataset 3D-MoIT. It is designed to cultivate 3D-MoLM’s ability to follow instructions, and to enhance its perception of 3D-dependent molecule properties. Specifically, 3D-MoIT is sourced from two databases: 1) PubChem, which offers a wide range of molecular properties, origins, and applications, and 2) PubChemQC (Nakata, 2015), which specializes in 3D-dependent molecular properties. As shown in Figure 3, for the PubChem portion, we leverage GPT-3.5 to generate QA pairs based on their descriptions. Yet, molecular properties collected from PubChem (e.g., molecular weight and LogP) can be largely inferred from 1D or 2D molecular data. To enhance 3D-MoIT’s perception of 3D molecular structures, we further incorporate data from PubChemQC, which includes 3D-dependent molecule properties (e.g., HOMO and LUMO; McQuarrie & Simon (1997)). We fill these properties into a set of text templates, transforming them into instruction tuning formats, as Figure 1 illustrates.

Our contributions can be summarized as follows:

- We propose 3D-MoLM, a new framework for 3D molecule-text interpretation. 3D-MoLM employs a 3D molecule-text projector to bridge the modality gap between a 3D molecular encoder and an LM, enabling the LM to perceive 3D molecular structures.
- We curate 3D-MoIT, a 3D molecule-centric instruction tuning dataset. We extract and transform data from PubChem and PubChemQC to an instruction following format, to cultivate 3D-MoLM’s ability in instruction following and 3D molecule-text interpretation.
- 3D-MoLM achieves state-of-the-art performances in extensive downstream tasks. Notably, on the PubChem Dataset, for molecule-text retrieval and molecule captioning, it outperforms baselines by



(a) Stage 1. The 3D molecule-text projector (*i.e.*, Q-Former) with the attached frozen 3D molecule encoder is optimized for 3D molecule-text representation learning. Stage 1 involves three training objectives: molecule-text matching, molecule-text contrasting, and molecule captioning.



(b) Stage 2 & 3. 3D-MoLM is trained to perform 3D molecule-to-text generations given 3D molecular tokens (extracted by the Q-former) and 1D textual prompt tokens.

Figure 2: Illustration of 3D-MoLM’s architectures at different stages.

20% accuracy and 6.47 ROUGE-L, respectively. Further, it surpasses the baselines with 1D or 2D molecular perceptions on open-text QA tasks, especially on 3D-dependent properties, verifying the capability of 3D molecule-text interpretation.

2 3D-MoLM: 3D MOLECULAR LANGUAGE MODELING

3D-MoLM incorporates a 3D molecular encoder into an LM, aiming to align 3D molecular geometries with textual concepts and facilitate a comprehensive cross-modal understanding of molecules. Consequently, 3D-MoLM is able to read 3D molecular structures, amplifying its molecular understanding and facilitating 3D-text interpretation. Our idea draws from related works in molecule-text modeling, multi-modal instruction tuning, and multi-modal LMs. See Appendix A for a comprehensive literature review. Here we delve into 3D-MoLM’s architecture and its training pipeline.

2.1 MODEL ARCHITECTURE

3D-MoLM’s architecture consists of three key components: 1) a 3D molecular encoder, focusing on encoding 3D molecular structures; 2) a 3D molecule-text projector, aiming to map the 3D molecular encoder’s representations to the input space of the LM; and 3) an LM, which specializes in text generation and is later adapted for understanding 3D molecular structures.

3D Molecular Encoder. We adopt Uni-Mol (Zhou et al., 2023) as our 3D molecular encoder f_{mol} . Specifically, Uni-Mol is pretrained on a large molecule dataset comprising 209M 3D molecular conformations. Formally, let $m = (\mathcal{V}, \mathbf{h}, \mathbf{C})$ denote a molecule, where \mathcal{V} and \mathbf{h} separately represent atomic nodes and their features, and $\mathbf{C} \in R^{|\mathcal{V}| \times 3}$ collects the 3D coordinates of nodes. In Uni-Mol, the representation for each pair of atoms is initialized using invariant spatial positional encoding derived from 3D coordinates \mathbf{C} . This encoding, grounded in the pair-wise Euclidean distances between atoms, ensures that the representation remains consistent regardless of global rotations or translations. Subsequently, representations of atoms and atom pairs engage in a self-attention mechanism, generating the molecular representation with 3D spatial information. Overall, the 3D molecular encoder f_{mol} performs molecule encoding procedure to obtain the atomic representations:

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{|\mathcal{V}|}] = f_{\text{mol}}(m), \quad (1)$$

where \mathbf{x}_i corresponds to the representation of the i -th atom.

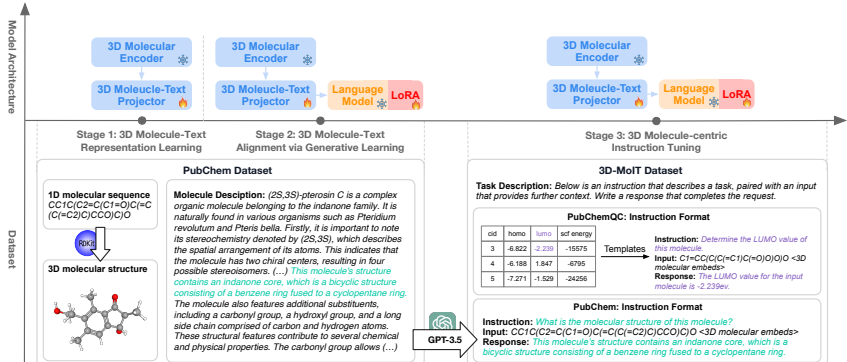


Figure 3: Illustration of the model architectures (upper part) and the dataset usage (bottom part) for the three training stages. PubChem is used for the stage 1 (i.e., 3D molecule-text representation learning) and stage 2 (i.e., 3D molecule-text alignment via generative learning). 3D-MoIT is used for 3D molecule-centric instruction tuning. Texts in the same color indicate the same information source.

3D Molecule-Text Projector. Taking inspiration from the leading vision-language models (Li et al., 2023; Dai et al., 2023), we architect the 3D molecule-text projector f_{pro} as a Querying Transformer (i.e., Q-Former) and initialize it from the Sci-BERT’s checkpoint (Beltagy et al., 2019). As illustrated in Figure 2a, Q-Former has two transformers with shared self-attention layers: one molecule transformer for processing 3D molecule features, and one text transformer for processing texts. The text transformer follows the same architecture of BERT (Devlin et al., 2019), while the molecule transformer adds cross-attention modules between the modules of self-attention and feed-forward to extract molecule features. Specifically, the molecule transformer maintains K learnable query tokens. Given 3D molecule input, the query tokens can interact with the 3D molecular encoder’s representations through the cross-attention modules. Therefore, the K query tokens’ output representations contain molecule information, represented as $\mathbf{M} = [m_1, m_2, \dots, m_K]$. The 3D molecule-text projector’s forward function can be written as:

$$\mathbf{M} = [m_1, m_2, \dots, m_K] = f_{\text{pro}}(\mathbf{X}). \quad (2)$$

Language Model (LM). We employ Llama2 (Touvron et al., 2023b) as our base LM f_{lm} to leverage its powerful text generation capability and internal chemistry knowledge. Although pretrained for general-purpose usage, the extensive biomedical literature in Llama 2’s pretraining corpus enables it to efficiently interpret 1D molecular sequences (e.g., SMILES) and proficiently address essential QA tasks that are relevant to molecular understanding. In this work, we let Llama2 process mixed token sequences that includes both textual tokens and 3D molecular tokens, which is detailed in Section 2.2.1. Formally, we denote a mixed token sequence that include l textual and molecular tokens as $\mathbf{Z} = [z_1, z_2, \dots, z_l]$. Further, the LM adopts a causal mask to generate textual response $\hat{\mathbf{Z}}$ with length n , where the prediction for the i -th token, \hat{z}_i , is dependent on its previous tokens:

$$\hat{\mathbf{Z}} = [\hat{z}_{l+1}, \hat{z}_{l+2}, \dots, \hat{z}_{l+n}], \quad \hat{z}_i = f_{\text{lm}}(\mathbf{Z}_{<i}), \quad \mathbf{Z}_{<i} = [z_1, z_2, \dots, z_l, \hat{z}_{l+1}, \dots, \hat{z}_{i-1}], \quad (3)$$

where each \hat{z}_i is later transformed by a linear layer f_{vocab} accompanied by a softmax function, converting into a probabilistic distribution over the vocabulary. The final prediction \tilde{z}_i for the i -th token is the word in the vocabulary with the highest probability, defined as:

$$\tilde{z}_i = \arg \max_{w \in \text{vocab}} f_{\text{vocab}}(\hat{z}_i)[w]. \quad (4)$$

2.2 MODEL TRAINING

To tackle the identified two challenges of 3D molecule-text alignment and 3D molecule-centric instruction tuning, we delineate a three-stage training pipeline (cf. Figure 3) for 3D-MoLM, including 1) 3D molecule-text representation learning, 2) 3D molecule-text alignment via generative learning, and 3) 3D molecule-centric instruction tuning.

2.2.1 3D MOLECULE-TEXT ALIGNMENT

Data Preparation – PubChem. A substantial collection of 3D molecule-text pairs is essential to bridge the gap between the 3D molecular representation space and the LM’s input space. We begin

Table 1: Statistics of 3D-MoIT. Comp.→computed property; Desc.→descriptive property.

Subset	PubChemQC		PubChem		
	#Mol	#Comp. QA	#Mol	#Comp. QA	#Desc. QA
Pretrain	3,119,717	12,478,868	301,658	1,199,066	1,508,290
Train	623,944	2,495,776	12,000	46,680	60,000
Valid	77,993	311,972	1,000	3,898	5,000
Test	77,993	311,972	2,000	7,785	10,000

by collecting molecular SMILES-text pairs from PubChem and further employ GPT-3.5 to enrich the less annotated molecular descriptions. Molecular 3D conformations are obtained by running the MMFF algorithm in RDKit (Landrum et al., 2013). As a result, we obtain 316K 3D molecule-text pairs for the alignment pretraining and downstream tasks of molecule-text retrieval and molecule captioning. The details of processing, subset split, and quality evaluation are in Appendix B.

Stage 1: 3D Molecule-Text Representation Learning. In the first stage, we jointly pretrain the Q-Former together with the frozen 3D molecular encoder on the collected 3D molecule-text pairs. Following BLIP-2 (Li et al., 2023), we perform multi-objective training, including molecule-text matching, molecule-text contrasting and molecule captioning (cf. Figure 2a). These training objectives are designed to cultivate Q-Former’s ability to extract molecular features that resonate profoundly with the corresponding text. Specifically, the molecule-text matching task mandates the model to differentiate between matched and unmatched pairs, enhancing the fine-grained alignment of cross-modal representations. In molecule-text contrasting, the similarity between a molecule and its corresponding text (*i.e.*, positive pair) is contrasted against those of negative pairs, aiming to maximize the mutual information between the molecular and textual representations. In molecule captioning, the Q-Former is trained to generate text descriptions, based on given molecular inputs.

Stage 2: 3D Molecule-Text Alignment via Generative Learning. In this stage, we connect the 3D molecular encoder and the Q-Former with the LM for 3D molecule-to-text generation, as illustrated in Figure 2b. The training objective is conditional language modeling: 3D-MoLM is trained to generate textual responses given 3D molecular tokens and 1D textual prompt tokens. Specifically, the 3D molecular tokens are extracted by the Q-Former, while the 1D textual prompt tokens include the molecule’s SMILES sequence and a textual task description. See Appendix E for ablation studies on prompt templates. This training objective encourages the LM to discern the contextual interplay between textual and molecular tokens, thus aiding in 3D molecule-text interpretation. For efficiency consideration, we freeze the 3D molecular encoder and employ LoRA tuning for the LM (Hu et al., 2021). Overall, this stage aligns 3D molecular representations with the LM’s textual input space, aiding in effective molecule-to-text generation.

2.2.2 3D MOLECULE-CENTRIC INSTRUCTION TUNING

Stage 3: 3D Molecule-centric Instruction Tuning. In the final stage, we freeze the 3D molecular encoder and conduct instruction fine-tuning to jointly optimize the 3D molecule-text projector and the LM. This fine-tuning has two purposes: 1) to enhance our model’s ability to follow various instructions; and 2) to improve our model’s understanding of 3D molecular structures, especially in recognizing 3D-dependent properties. It is framed as conditional text generation optimization based on the input prompt and the standard language modeling loss is adopted. Now, we elaborate on the details of our instruction tuning dataset 3D-MoIT.

Data Preparation – 3D-MoIT. 3D-MoIT sources data from the PubChem (Kim et al., 2021) and PubChemQC (Nakata, 2015) databases. Specifically, the instruction tuning data from PubChem can be divided into two types: computed molecular properties and descriptive texts, where computed molecular properties are numerical values and descriptive texts characterize molecule properties in free texts. To effectively activate the 3D perception, we further include computed molecular properties from PubChemQC. PubChemQC contains 3D-dependent molecule attributes, which mostly cannot be inferred from 1D or 2D molecular representations. PubChemQC also includes DFT-determined 3D conformations, commonly regarded as the ground truth in 3D molecular modeling. The statistics of 3D-MoIT are shown in Table 1. We detail the dataset construction as follows:

- **PubChem: Instruction Format.** We select the following computed properties: molecular weight, LogP, TPSA, and complexity. They are transformed into instruction tuning format using a pre-defined set of text templates. For descriptive text, we adopt GPT-3.5 to read molecular descriptions

Table 2: Molecule-Text retrieval results on the PubChem Dataset. † denotes pretraining on the original PubChem texts without GPT-3.5 enrichment. We report performances of both using molecule to retrieve text (M2T) and using text to retrieve molecule (T2M).

Model	Retrieval in batch				Retrieval in test set			
	M2T (%)		T2M (%)		M2T (%)		T2M (%)	
	Acc	R@20	Acc	R@20	Acc	R@20	Acc	R@20
1D SMILES								
Sci-BERT	85.32	98.74	84.20	98.43	41.67	87.31	40.18	86.77
KV-PLM	86.05	98.63	85.21	98.47	42.80	88.46	41.67	87.80
2D Graph								
MoMu-S	87.58	99.24	86.44	99.38	47.29	90.77	48.13	89.92
MoMu-K	88.23	99.41	87.29	99.42	48.47	91.64	49.46	90.73
MoMu-S [†]	90.43	99.53	89.38	99.60	60.51	93.24	58.36	91.35
MoMu-K [†]	90.89	99.67	90.16	<u>99.44</u>	62.07	93.06	59.17	92.01
3D Conformation								
3D-MoLM [†]	94.48	<u>99.74</u>	94.78	99.34	72.06	96.42	71.30	95.96
3D-MoLM	<u>93.50</u>	100.00	<u>92.89</u>	99.59	<u>69.05</u>	<u>95.91</u>	<u>70.13</u>	<u>94.88</u>

and generate five QA pairs for each molecule, as illustrated in Figure 3. PubChem includes a diverse range of molecule properties, enhancing the comprehensiveness of 3D-MoIT.

- **PubChemQC: Instruction Format.** We select the following computed molecular properties: HOMO, LUMO, HOMO-LUMO Gap, and SCF-energy. These properties are transformed into instruction tuning format by filling the values into pre-defined text templates, as illustrated in Figure 3. We use the processed dataset released by (Xu et al., 2021) and follow the scaffold split.

3 EXPERIMENT

In this section, we conduct extensive experiments, including molecule-text retrieval, molecule captioning, and open-text molecular QA tasks, to demonstrate the effectiveness of 3D-MoLM for 3D molecule-text interpretation. See Appendix C for experimental details of each downstream task.

3.1 MOLECULE-TEXT RETRIEVAL

We assess the Stage-1 checkpoint of 3D-MoLM on the downstream subsets of PubChem Dataset for molecule-text retrieval. These subsets contain real-world molecules paired with textual descriptions longer than 20 words. We opt not to evaluate existing molecule-text datasets of PCDes (Zeng et al., 2022), because of data leakage. PCDes and our curated PubChem dataset stem from the same source (*i.e.*, some PCDes test data present in our pretraining set). We employ Sci-BERT (Beltagy et al., 2019), KV-PLM (Zeng et al., 2022), and MoMu (Su et al., 2022) as baselines and evaluate the performance by Accuracy and Recall@20, both within a batch of 64 samples and across the entire test set. Baselines are initiated from their official checkpoints and finetuned using the downstream partition of the PubChem Dataset, except MoMu[†], which is our re-implementation with the original PubChem texts without GPT-3.5 enrichment. From Table 2, we have the following observations:

1) 3D-MoLM surpasses existing baselines, including both molecular 1D-language models (*i.e.*, Sci-BERT, KV-PLM) and 2D-language models (*i.e.*, MoMu-S/K), by a large margin. The improvement can be attributed to two pivotal factors. Firstly, the performance gain 3D-MoLM[†] over MoMu[†], which are both pretrained on the same PubChem molecule-text pairs, demonstrates that Q-Former benefits from multi-objective pretraining across diverse granularities, distinguishing it from other models that predominantly rely on the molecule-text contrastive learning objective. Secondly, the scale of the PubChem Dataset, which we curated, offers a considerable advantage, which is verified by the performance gain MoMu[†] over the original MoMu. With the collection of 301K molecule-text pairs for pretraining, it markedly surpasses the 15K pairs in the MoMu *w.r.t.* scale. This increased scale provides our model with a richer and more diverse set of training examples, allowing it to better capture the molecular-text interactions. Such superior performance underscores the effectiveness of 3D molecular representation learning in Stage 1, demonstrating the capability of Q-Former to extract molecular features that are strongly related to the textual descriptions.

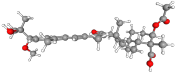
2) The retrieval performance on the PubChem test set appears to be negatively impacted by GPT-3.5 enrichment. We infer that this decline is caused by the enrichment process enlarging the distribution gap between the pretraining and downstream subsets. While the original texts might be succinct and

Table 3: Molecule captioning results on PubChem Dataset. † denotes pretraining on the original PubChem texts without GPT-3.5 enrichment. Llama2-7B, without a molecule-text projector, goes through Stage 2 training with the prompt of 1D SMILES. 2D-MoLM replaces the 3D molecular encoder with a 2D molecular encoder and goes through the same training process as 3D-MoLM.

(a) Molecule captioning results.

Type	Model	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	METEOR	
	1D SMILES							
	MolT5-Small	22.53	15.23	30.44	13.45	20.30	23.98	
	MolT5-Base	24.51	16.61	32.19	14.04	21.35	26.10	
	MolT5-Large	25.87	17.28	34.07	16.42	23.41	28.04	
Specialist	1D SMILES + 2D Graph							
	MoMu-Small	22.86	16.01	30.98	13.65	20.75	24.35	
	MoMu-Base	24.74	16.77	32.45	14.62	22.09	27.16	
	MoMu-Large	26.34	18.01	34.75	16.86	24.76	28.73	
	1D SMILES + 3D Conformation							
	3D-MoLM†	<u>29.82</u>	<u>22.39</u>	37.23	<u>22.49</u>	<u>31.07</u>	<u>32.69</u>	
	3D-MoLM	30.32	22.52	<u>36.84</u>	22.32	31.23	33.06	
Generalist	1D SMILES							
	Llama2-7B	27.01	20.94	35.76	20.68	28.88	32.11	
	1D SMILES + 2D Graph							
	2D-MoLM	27.15	21.19	36.02	20.76	29.12	32.28	
		1D SMILES + 3D Conformation						
		3D-MoLM†	29.25	22.07	<u>36.48</u>	21.80	30.95	<u>33.12</u>
	3D-MoLM	<u>28.95</u>	<u>21.63</u>	36.51	<u>21.26</u>	<u>30.02</u>	33.55	

(b) A sample of molecule captioning.

Molecule	Ground Truth	3D-MoLM
	Globostellatic acid C is a tricyclic triterpenoid of the isomalabaricane group. It has a role as an anti-neoplastic agent and a metabolite. It is a tricyclic triterpenoid, an acetate ester, an enone, an ether and an oxo monocarboxylic acid . It is a conjugate acid of a globostellatate C(1-).	Globostellatic acid B is a tricyclic triterpenoid with formula C34H48O6, originally isolated from Globostellata. It has a role as a metabolite. It is a tricyclic triterpenoid, a monocarboxylic acid and an acetate ester . Globostellatic acid B is a natural product found in Globostellata with data available.

hence less informative, they exhibit word patterns more congruent with those in the test set, in stark contrast to the distinctive GPT-style exhibited by the enriched texts. Nonetheless, we argue that this enrichment process benefits forging a more comprehensive connection between molecular structures and their inherent properties. This assertion is further substantiated by subsequent experiments of textual generation tasks, where GPT-3.5 enrichment boosts the performance of 3D-MoLM.

3.2 MOLECULE CAPTIONING

Generalist and Specialist Models. For generation tasks of molecule captioning and open-text molecule QA, we present 3D-MoLM’s performances of two different variants: the specialist, which is fine-tuned for each generation task separately, and the generalist, which is a unified model trained on all the generation tasks. Both models are obtained by fine-tuning the Stage-2 checkpoint.

Here we evaluate 3D-MoLM for molecule captioning on the PubChem Dataset. We opt not to use CheBI-20 (Edwards et al., 2022) for two reasons: firstly, its origin from PubChem poses potential data leakage issues, akin to the concerns with PCDes; secondly, CheBI-20 is curated in a way that molecular names are replaced with “the molecule”, driving the model to emphasize properties instead of names. However, the molecular nomenclature system is intrinsically related to the identification of distinct molecular structures, encompassing features such as hydrocarbon chains and benzene rings. Consequently, a model’s ability to accurately predict these names serves as a testament to its adeptness in comprehending the underlying molecular structures. Thus, we elevate the task’s complexity by retaining molecular names in the texts, positioning this task as a combination of molecule captioning without molecular names (Edwards et al., 2022), and name prediction (Favre & Powell, 2013). To demonstrate the effectiveness of 3D molecular perception, we include 3D-MoLM’s variants of 1D (i.e., Llama2-7B) and 2D (i.e., 2D-MoLM) perception as baselines. Specifically, Llama2-7B, without a molecule-text projector, goes through Stage 2 & 3 training using 1D SMILES as molecule representations. 2D-MoLM replaces the 3D molecular encoder with a 2D molecular encoder (Liu et al., 2022b), and undergoes the same training process as 3D-MoLM. Specialist models

are fine-tuned using the training set from the PubChem Dataset. Table 3 presents the performances with metrics of BLEU, ROUGE, and METEOR, accompanied by a concrete sample. We observe that:

- 1) 3D-MoLM demonstrates superior performance across the board, with the highest scores on all evaluation metrics. While slightly erring in identifying the molecule as “Globostellatic acid B” rather than “Globostellatic acid C”, it pinpoints roles and primary structures such as the tricyclic triterpenoid, acetate ester, and monocarboxylic acid. This demonstrates the effectiveness of 3D molecule-text alignment training to bridge the gap between 3D molecular representations and LM’s input space. We also provide detailed analysis and discussion on failure cases in Appendix D.
- 2) The enrichment process via GPT bolsters the text generation capability based on 3D molecular structures. This underscores the hypothesis that the enrichment strengthens the connection between molecular structures and their associated properties, enhancing cross-modal molecular understanding. Similar gain can be observed in the following open-text QA tasks as well.
- 3) 3D-MoLM’s heightened performance, when juxtaposed with finetuned Llama2-7B and 2D-MoLM subjected to a similar training process but modeling molecules as 1D SMILES and 2D graphs, highlights the pivotal role of 3D structure perception in bolstering molecular understanding.

3.3 OPEN-TEXT MOLECULAR QUESTION-ANSWERING (QA)

We evaluate 3D-MoLM for open-text molecular QA on the 3D-MoIT dataset. Considering that open-text molecular QA is mostly unexplored in existing works, we mainly compare 3D-MoLM with its variants of 1D or 2D molecular perceptions. Notably, we report performances of specialists, which are trained for each task separately, and generalists, which are unified models trained on all the generation tasks. Table 4 presents the quantitative evaluations and QA samples. We observe that:

Observations for Descriptive Property QA: 1) Superiority of 3D-MoLM over baselines. It exhibits a commendable performance in 3D molecule-text interpretation, clearly surpassing the baselines. Specifically, It correctly identifies beryllium acetate as a solid that is soluble in water and goes beyond the ground truth by providing supplementary details, such as pinpointing the elevated melting and boiling points and attributing them to the potent intermolecular forces between beryllium and acetate, which underlines its depth of molecular understanding.

2) Benefit of descriptive instructions. Incorporating them amplifies the model’s molecular comprehension. This can be verified by that through instruction tuning, Llama2-7B (generalist) advances over its initialization point, manifesting improvements of 2.46 in BLEU-2 and 3.22 in METEOR.

Observations for Computed Property QA: 1) 3D-MoLM achieves superior performances on computed property QA task, consistently achieving the lowest MAE, especially on those properties intrinsically determined by 3D conformations (*i.e.*, highlighted properties in Table 4c). A remarkable performance lead of 0.77 eV among generalists on HOMO-LUMO Gap accentuates 3D-MoLM’s adeptness at 3D molecular understanding, which we attribute to its 3D perception. However, for properties that mainly stem from atomic compositions and interatomic connections (*i.e.*, molecule weight, LogP, TPSA, and complexity), the advantage, while persistent, is more subdued. This aligns with Uni-Mol, which displays larger advantages over 2D molecular models for predicting quantum mechanical properties grounded in 3D structures.

2) Efficacy of instruction tuning. Instruction tuning amplifies the model’s capability to adhere to and act upon instructions. This is illuminated by Llama2-7B’s official checkpoint’s occasional ineptitude in generating valid numerical responses, particularly when tasked with approximating 3D-dependent computed properties that aren’t straightforwardly derivable from SMILES.

3) Comparison with Uni-Mol. 3D-MoLM can enhance the accuracy of molecular property prediction by leveraging both the rich contextual knowledge found in chemistry literature and 3D molecular conformations. For instance, the pretraining corpus of 3D-MoLM contains descriptions of hydrophobicity (LogP) and solubility (TPSA). While Uni-Mol excels at predicting molecular properties by interpreting 3D conformations, it cannot utilize textual descriptions of chemical properties. This dual-source knowledge utilization can potentially enhance the prediction of molecular properties.


Observations for generalist and specialist: While the generalist model slightly underperforms in comparison to the specialist models, it still exhibits a performance gain over other baselines. This underscores 3D-MoLM’s versatility and capability to effectively handle multiple tasks.

Table 4: Open-text QA results on 3D-MoIT. * denotes the official checkpoint without any finetuning. † denotes molecule-text alignment on the original PubChem texts without GPT-3.5 enrichment. Llama2-7B, without a molecule-text projector, goes through Stage 3 instruction tuning by modeling molecules as 1D SMILES. 2D-MoLM replaces the 3D molecular encoder with a 2D molecular encoder and goes through three-stage training as 3D-MoLM.

(a) Descriptive property QA results.

Type	Model	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	METEOR
Specialist	Llama2-7B	28.15	23.24	35.14	22.08	30.41	46.87
	2D-MoLM	<u>30.84</u>	<u>25.09</u>	38.46	<u>24.22</u>	<u>33.04</u>	50.92
	3D-MoLM†	30.33	24.47	<u>38.48</u>	23.93	32.98	<u>51.33</u>
	3D-MoLM	32.00	26.13	40.13	25.55	34.64	52.15
Generalist	Llama2-7B*	25.22	21.16	31.48	19.21	25.22	43.17
	Llama2-7B	27.68	22.81	34.73	21.55	29.91	46.39
	2D-MoLM	<u>30.23</u>	<u>24.57</u>	37.85	22.95	32.28	50.08
	3D-MoLM†	29.92	24.44	<u>38.62</u>	<u>22.83</u>	<u>32.30</u>	50.81
	3D-MoLM	31.81	26.08	40.13	25.87	34.99	51.93

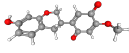
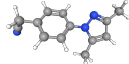
(b) A concrete sample of descriptive property QA.

Molecule	Question	Ground truth	3D-MoLM
	What are the physical properties of beryllium acetate?	Beryllium acetate is typically a white or colorless solid and is soluble in water .	On the physical properties of beryllium acetate, it is a white crystalline solid that is soluble in water . It has a relatively high melting point and boiling point due to the strong intermolecular forces between the beryllium and acetate.

(c) Computed property QA results. We report the MAE with a valid answer rate (%) – LMs sometimes fail to follow instructions to generate numerical responses. We highlight properties largely dependent on 3D structures.

Type	Model	Weight (g/mol)	LogP	TPSA (Å ²)	Complexity	HOMO (eV)	LUMO (eV)	H-L Gap (eV)	SCF (10 ⁴ eV)
Non-LM	Uni-Mol	20.35	0.59	13.48	57.24	0.32	0.35	0.21	0.45
Specialist	Llama2-7B	22.10 (96%)	1.45 (95%)	15.87 (92%)	69.74 (93%)	1.24 (96%)	1.04 (95%)	0.88 (92%)	0.70 (99%)
	2D-MoLM	21.48 (94%)	<u>0.88 (96%)</u>	13.52 (92%)	55.74 (94%)	0.92 (98%)	0.80 (96%)	0.67 (93%)	0.71 (99%)
	3D-MoLM†	<u>16.18 (96%)</u>	0.95 (96%)	<u>10.26 (94%)</u>	<u>49.15 (95%)</u>	<u>0.45 (98%)</u>	<u>0.36 (96%)</u>	<u>0.41 (94%)</u>	<u>0.39 (99%)</u>
	3D-MoLM	14.79 (95%)	0.66 (97%)	9.71 (93%)	44.85 (94%)	0.26 (97%)	0.25 (94%)	0.28 (94%)	0.35 (99%)
Generalist	Llama2-7B*	42.18 (82%)	2.10 (85%)	27.11 (84%)	121.87 (76%)	2.87 (70%)	1.89 (71%)	1.86 (70%)	3.84 (23%)
	Llama2-7B	27.42 (92%)	1.78 (93%)	17.07 (90%)	78.16 (92%)	1.89 (90%)	1.26 (90%)	1.25 (91%)	0.87 (99%)
	2D-MoLM	20.80 (92%)	1.36 (94%)	12.47 (89%)	<u>52.70 (91%)</u>	1.52 (93%)	1.13 (92%)	1.09 (88%)	0.96 (99%)
	3D-MoLM†	19.54 (93%)	<u>0.92 (92%)</u>	11.14 (92%)	54.68 (90%)	0.65 (94%)	<u>0.41 (92%)</u>	<u>0.55 (89%)</u>	<u>0.49 (99%)</u>
	3D-MoLM	16.58 (92%)	0.78 (95%)	10.90 (90%)	45.49 (89%)	0.35 (95%)	0.36 (93%)	0.32 (90%)	0.38 (98%)

(d) Concrete samples of computed property QA.

Molecule	Question	Ground truth	3D-MoLM
	Please provide the Molecular Weight value for this molecule. If uncertain, provide an estimate. Respond with the numerical value only.	The Molecular Weight for the input molecule is 286.28 g/mol .	Input molecule has a Molecular Weight of 288.30 g/mol .
	Could you give me the HOMO-LUMO Gap value of this molecule? If uncertain, provide an estimate. Respond with the numerical value only.	The HOMO-LUMO Gap for the input molecule is 5.325 eV .	The HOMO-LUMO Gap for the input molecule is 5.762 eV .

4 CONCLUSION

In this work, we introduce 3D-MoLM, a new approach tailored for 3D-molecular language modeling. 3D-MoLM equips an LM with a 3D molecular encoder for 3D molecule-text interpretation. This is achieved by a 3D molecule-text projector that aims to map 3D molecular representations into the LM’s textual space. Furthermore, 3D-MoLM incorporates 3D molecule-centric instruction tuning, enhancing both its adherence to human instructions and 3D molecular understanding. Extensive experiments demonstrate that 3D-MoLM excels in various tasks, including molecule-text retrieval, molecule captioning, and open-text molecular QA.

Despite the promising results, our work has a few limitations. Compared with vision-language modeling methods (Dai et al., 2023; Hong et al., 2023), the scale of our 3D molecule-text dataset is notably constrained, inherently limiting the performance. This motivates our search for high-quality texts closely related to 3D molecular structures. Furthermore, this study does not explore other intriguing capabilities of large LMs, such as in-context learning and chain-of-thought reasoning.

5 ACKNOWLEDGEMENTS

This research is supported by the National Natural Science Foundation of China (92270114) and partially supported by the National Research Foundation Singapore under the AI Singapore Programme (AISG Award No: AISG2-TC-2023-010-SGIL), the Singapore Ministry of Education Academic Research Fund Tier 1 (Award No: T1 251RES2207) and the Google Cloud Research Credits program (Award No: 6NW8-CF7K-3AG4-1WH1). This research is also supported by the CCCD Key Lab of Ministry of Culture and Tourism and NExT Research Center.

REFERENCES

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L. Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022.
- Iz Beltagy, Kyle Lo, and Arman Cohan. Scibert: A pretrained language model for scientific text. In *EMNLP/IJCNLP (1)*, pp. 3613–3618. Association for Computational Linguistics, 2019.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3558–3568, 2021.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models. *CoRR*, abs/2210.11416, 2022.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *NeurIPS*, 2023.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pp. 4171–4186. Association for Computational Linguistics, 2019.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model. In *ICML*, 2023.
- Carl Edwards, ChengXiang Zhai, and Heng Ji. Text2mol: Cross-modal molecule retrieval with natural language queries. In *EMNLP (1)*, pp. 595–607. Association for Computational Linguistics, 2021.
- Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. Translation between molecules and natural language. In *EMNLP*, pp. 375–413. Association for Computational Linguistics, 2022.
- Junfeng Fang, Wei Liu, Yuan Gao, Zemin Liu, An Zhang, Xiang Wang, and Xiangnan He. Evaluating post-hoc explanations for graph neural networks via robustness analysis. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a.
- Junfeng Fang, Xinglin Li, Yongduo Sui, Yuan Gao, Guibin Zhang, Kun Wang, Xiang Wang, and Xiangnan He. Exgc: Bridging efficiency and explainability in graph condensation. In *WWW*. ACM, 2024a.

- Junfeng Fang, Shuai Zhang, Chang Wu, Zhiyuan Liu, Sihang Li, Kun Wang, Wenjie Du, Xiang Wang, and Xiangnan He. Molc: Towards molecular relational modeling in language models, 2024b.
- Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang, Hua Wu, and Haifeng Wang. Geometry-enhanced molecular representation learning for property prediction. *Nat. Mach. Intell.*, 4(2):127–134, 2022.
- Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Huajun Chen. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. *arXiv preprint arXiv:2306.08018*, 2023b.
- Henri A Favre and Warren H Powell. *Nomenclature of organic chemistry: IUPAC recommendations and preferred names 2013*. Royal Society of Chemistry, 2013.
- Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *arXiv preprint arXiv:2307.12981*, 2023.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *CoRR*, abs/2106.09685, 2021.
- Robert O Jones. Density functional theory: Its origins, rise to prominence, and future. *Reviews of modern physics*, 87(3):897, 2015.
- William L Jorgensen. The many roles of computation in drug discovery. *Science*, 303(5665):1813–1818, 2004.
- Martin Karplus and J Andrew McCammon. Molecular dynamics simulations of biomolecules. *Nature structural biology*, 9(9):646–652, 2002.
- Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A. Shoemaker, Paul A. Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan Bolton. Pubchem in 2021: new data content and improved web interfaces. *Nucleic Acids Res.*, 49(Database-Issue):D1388–D1395, 2021.
- Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alán Aspuru-Guzik. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach. Learn. Sci. Technol.*, 1(4):45024, 2020.
- Greg Landrum et al. Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum*, 8:31, 2013.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. *CoRR*, abs/2301.12597, 2023.
- Sihang Li, Xiang Wang, An Zhang, Yingxin Wu, Xiangnan He, and Tat-Seng Chua. Let invariant rationale discovery inspire graph contrastive learning. In *International conference on machine learning*, pp. 13052–13065. PMLR, 2022a.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL/IJCNLP (1)*, pp. 4582–4597. Association for Computational Linguistics, 2021.
- Yanguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *ICLR*. OpenReview.net, 2022b.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023a.
- Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Anima Anandkumar. Multi-modal molecule structure-text model for text-based retrieval and editing. *CoRR*, abs/2212.10789, 2022a.
- Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, and Jian Tang. Pre-training molecular graph representation with 3d geometry. In *ICLR*. OpenReview.net, 2022b.

- Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. Molca: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In *EMNLP*, 2023b. URL <https://openreview.net/forum?id=14WRhMNq7H>.
- Zhiyuan Liu, Yaorui Shi, An Zhang, Enzhi Zhang, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. Rethinking tokenizer and decoder in masked graph modeling for molecules. In *NeurIPS*, 2023c. URL <https://openreview.net/forum?id=fWLF8DV0fI>.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688*, 2023.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- Shuqi Lu, Zhifeng Gao, Di He, Linfeng Zhang, and Guolin Ke. Highly accurate quantum chemical property prediction with uni-mol+. *CoRR*, abs/2303.16982, 2023.
- Chenyang Lyu, Minghao Wu, Longyue Wang, Xinting Huang, Bingshuai Liu, Zefeng Du, Shuming Shi, and Zhaopeng Tu. Macaw-llm: Multi-modal language modeling with image, audio, video, and text integration. *arXiv*, 2023.
- Donald Allan McQuarrie and John Douglas Simon. *Physical chemistry: a molecular approach*, volume 1. University science books Sausalito, CA, 1997.
- Maho Nakata. The pubchemqc project: a large chemical database from the first principle calculations, 2015.
- OpenAI. Chatgpt, 2023a. URL <https://openai.com/blog/chatgpt>.
- OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023b.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- Yaorui Shi, An Zhang, Enzhi Zhang, Zhiyuan Liu, and Xiang Wang. ReLM: Leveraging language models for enhanced chemical reaction prediction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 5506–5520, December 2023.
- Bing Su, Dazhao Du, Zhao Yang, Yujie Zhou, Jiangmeng Li, Anyi Rao, Hao Sun, Zhiwu Lu, and Ji-Rong Wen. A molecular multimodal foundation model associating molecule graphs with natural language. *CoRR*, abs/2209.05481, 2022.
- Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. Galactica: A large language model for science. *CoRR*, abs/2211.09085, 2022.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. In *NeurIPS*, pp. 200–212, 2021.

- Alain C Vaucher, Philippe Schwaller, Joppe Geluykens, Vishnu H Nair, Anna Iuliano, and Teodoro Laino. Inferring experimental procedures from text-based representations of chemical reactions. *Nature communications*, 12(1):2573, 2021.
- Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023.
- Zhao Xu, Youzhi Luo, Xuan Zhang, Xinyi Xu, Yaochen Xie, Meng Liu, Kaleb Dickerson, Cheng Deng, Maho Nakata, and Shuiwang Ji. Molecule3d: A benchmark for predicting 3d geometries from molecular graphs, 2021.
- Zheni Zeng, Yuan Yao, Zhiyuan Liu, and Maosong Sun. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications*, 13(862), 2022.
- Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv*, 2023.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Min Wu, Tingjun Hou, and Mingli Song. Root-aligned smiles: a tight representation for chemical reaction prediction. *Chemical Science*, 13(31):9023–9034, 2022.
- Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. In *ICLR*. OpenReview.net, 2023.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

A RELATED WORKS

In this section, we provide a review of literature related to molecule-text modeling, multi-modal instruction tuning, and multi-modal language models.

Molecule-Text Modeling. Early research approaches molecular understanding by representing molecules as 1D sequences. Specifically, KV-PLM (Zeng et al., 2022) represents molecules using 1D SMILES and employs a masked-language-modeling loss for pretraining on biomedical texts. MolT5 (Edwards et al., 2022) is a T5-based (Raffel et al., 2020) model tailored for molecular tasks (*i.e.*, SMILES-to-text and text-to-SMILES translations). Mol-Instruct (Fang et al., 2023b) utilizes the 1D SELFIES (Krenn et al., 2020) molecular descriptor and offers an instruction dataset for biomolecular research. Subsequently, 2D molecular graph encoders (Liu et al., 2023c) are introduced to explicitly capture topological information. For example, MoMu (Su et al., 2022) and MoleculeSTM (Liu et al., 2022a) use cross-modal contrastive learning to bridge the representation spaces of molecular graphs and texts. MolCA (Liu et al., 2023b) combines 1D SMILES and 2D graph representations for molecule-to-text generation. ReLM (Shi et al., 2023) applies ChatGPT to rerank reaction prediction results proposed by a GNN model. However, current molecule-text modeling approaches mostly rely on 1D sequences or 2D graphs for molecular representation, overlooking the vital spatial information in 3D molecular conformations.

Multi-Modal Instruction Tuning. Instruction tuning (Chung et al., 2022; Longpre et al., 2023) trains LMs to adhere to natural language instructions, enhancing their ability to generalize to unseen tasks and instructions. It has been extended in the multi-modal domain to augment the LM’s cross-modal comprehension. For instance, InstructBLIP (Dai et al., 2023) adapts 26 vision-language datasets to the instruction tuning format and introduces instruction-aware visual feature extraction, allowing adaptive and informative feature extraction based on provided instructions. LLaVA (Liu et al., 2023a) and MiniGPT-4 (Zhu et al., 2023) leverage GPT-3.5/GPT-4 (OpenAI, 2023a;b) to convert or polish image-text pairs into high-quality instruction-following format, facilitating integrated visual and language understanding. In the 3D visual domain, 3D-LLM (Hong et al., 2023) curates a 3D-language instruction tuning dataset to instill a 3D perspective into LMs. They have showcased the potential of instruction tuning to enhance the cross-modal understanding of LMs. However, its application in the 3D molecule-text domain remains largely unexplored.

Multi-Modal Language Models. Multi-modal LMs aim to merge the realms of text with other modalities. Early vision-language models (Radford et al., 2021; Li et al., 2022b) employ contrastive learning to fuse the visual and textual representations. Recent studies (Li et al., 2023; Alayrac et al., 2022; Tsimpoukelli et al., 2021; Driess et al., 2023) suggest that visual feature spaces can be seamlessly aligned with the input spaces of LMs, enabling them to undertake language generation tasks based on visual inputs. Aside from images, researchers have integrated preceptors of diverse modalities, including video (Zhang et al., 2023), audio (Lyu et al., 2023) and 3D point clouds (Xu et al., 2023), into LMs, giving rise to a series of multi-modal LMs. We draw inspiration from these revelations to devise 3D-MoLM.

B PUBCHEM DATASET AND ENRICHMENT PROCESS

GPT-3.5 Enrichment. We begin our data preparation by collecting 324,689 molecular SMILES-text pairs from PubChem (Kim et al., 2021). However, they include a large number of uninformative texts – the shortest description has only 1 word and 82.29% of the collected molecules have descriptions less than 20 words. To ensure the effective evaluation for alignment, we first extract a high-quality subset of 15K pairs encompassing descriptions with more than 19 words. This curated subset is subsequently partitioned into train / validation / test sets containing 12K / 1K / 2K pairs, respectively. The less informative remainder is reserved for alignment purpose (*i.e.*, stage 1&2 pretraining). Yet, we argue that those less informative texts potentially offers an insufficient perspective on molecular properties and structures, which could compromise the 3D molecule-text alignment. To address this shortfall, as illustrated in Figure 4, we employ GPT-3.5 to enhance these molecule descriptions, thus bolstering the less annotated molecules. We note that the enrichment process using GPT is exclusively applied to the pretraining set, ensuring that evaluations in downstream tasks (*i.e.*, train / validation / test sets) are conducted on original real-world data.

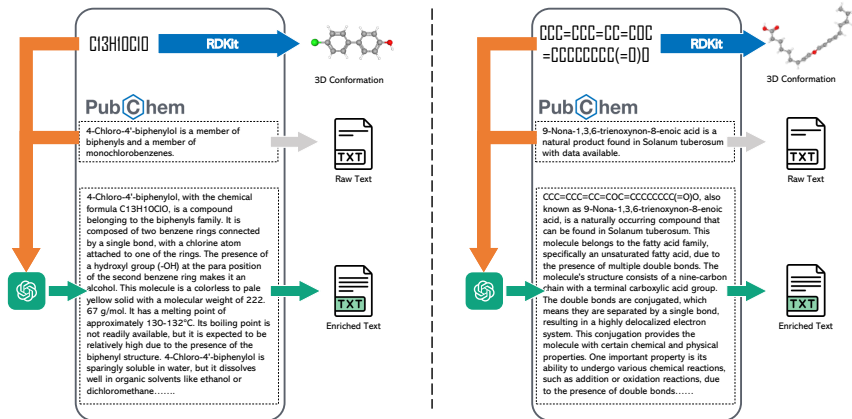


Figure 4: Utilizing GPT-3.5 to enrich the descriptions in the pretraining subset of PubChem Dataset.

Table 5: Number of molecule-text pairs in subsets of PubChem Dataset.

Subset	#Molecule-Text Pair	#Min Word	#Avg Word
Pretrain	301,658	1	17.84
Train	12,000	20	57.24
Valid	1,000	20	58.31
Test	2,000	20	55.21

Correctness and Comprehensiveness Evaluation. To evaluate the correctness and comprehensiveness of the original texts in PubChem and those enriched by GPT-3.5, we solicited evaluations from three human experts in chemistry and also from GPT-4 with the same score standards in Figure 5. These evaluations, conducted on subsets of 100 (for human experts) and 1,000 (for GPT-4) samples respectively, are presented in Figure 6. The results indicate that the enhanced descriptions have superior comprehensiveness, while maintaining high correctness. Adopting the GPT-3.5 generation and subsequent GPT-4 evaluation approach, rather than directly using GPT-4 for text enrichment, was primarily driven by cost considerations; utilizing GPT-4 for the enrichment would have incurred expenses exceeding \$5000.

Raw 3D conformation Generation. Subsequently, we leverage RDKit to transform the 1D SMILES into raw 3D conformations. While the raw conformations derived from 1D SMILES might not match the accuracy of costly DFT (Jones, 2015) equilibrium conformations, which necessitate several hours for single molecule property calculations, they suffice for preliminary cross-modal alignment. Meanwhile, RDKit occasionally fails to generate 3D conformations. As such, pairs failing this generation process are excluded. As shown in Table 5, the resultant dataset split for pretrain/train/validation/test comprises 301,658 / 12,000 / 1,000 / 2,000 3D molecule-text pairs, respectively. The pretraining set with GPT-enriched descriptions is utilized for 3D molecule-text alignment, detailed in Section 2.2.1. Meanwhile, the train, validation, and test sets are used for cross-modal retrieval and molecule captioning tasks, detailed in Section 3.1 and 3.2, serving as a measure of alignment quality.

C EXPERIMENT SETUPS

In this section, we present the experimental setups for three pretraining stages and downstream tasks.

Stage 1 Pretrain. The Q-former attached with a frozen 3D molecular encoder is pretrained for 50 epochs and the number of query tokens in it is set to 8. AdamW (Loshchilov & Hutter, 2018) optimizer is adopted with a weight decay of 0.05 and a learning rate scheduler of a combination of linear warmup with 1000 steps and cosine decay, in which the peak and minimal learning rates are $1e-4$ and $5e-6$, respectively. And the batch size and maximal text length are 64 and 256, respectively. The computation overhead is 40 GPU hours on NVIDIA A100 with BFloat16 Mixed precision.

Molecule-text retrieval. The Stage-1 checkpoint is finetuned for 10 epochs and evaluated on the downstream subsets (*i.e.*, train/validation/test) of PubChem with the same optimizer and learning rate scheduler configurations, with the exception of the warmup steps, which are set to 200, as stage 1

Description enrichment prompt template

Instruction: You are an expert chemist. Given the molecular SMILES and its raw description, your task is to polish the raw description to provide the correct and detailed description of the molecule's chemical and physical properties using your chemical molecular knowledge.

Input: Molecule SMILES: [SMILES]
Raw Description: [Raw text]

Output:

Correctness score standard for molecular description

Instruction: You are an expert chemist and I will give you the SMILES of a molecule and a paragraph of textual description for it. Your task is to evaluate the correctness of this textual description using your chemical molecular knowledge and score it from 1 to 5 according to following standard.

Molecule Description Correctness Score Standard (1-5)

- Inaccurate:**
 - Contains multiple factual errors or inconsistencies.
 - Makes use of non-standard or incorrect nomenclature.
 - Provides incorrect context regarding the molecule's properties or significance.
- Partially Correct:**
 - Contains one or more minor factual mistakes or inconsistencies.
 - Uses a mix of standard and non-standard nomenclature.
 - Provides limited context, which might be partly inaccurate.
- Generally Correct:**
 - Contains minor omissions, but no glaring factual errors.
 - Mostly uses standard nomenclature with minor deviations.
 - Gives a basic, correct overview of the molecule's significance or properties.
- Mostly Correct:**
 - Provides a description with most details being accurate.
 - Contains minute inaccuracies but nothing that significantly detracts from the overall correctness.
 - Consistently employs standard nomenclature with minor exceptions.
- Highly Correct:**
 - Offers a holistic and detailed description of the molecule without any discernible inaccuracies.
 - Perfectly adheres to standard nomenclature and conventions.
 - Furnishes a thorough understanding of the molecule's significance, backed by accurate context and information.

Input: SMILES: [SMILES]
Description: [Molecular text]

Output:

Comprehensiveness score standard for molecular description

Instruction: You are an expert chemist and I will give you the SMILES of a molecule and a paragraph of textual description for it. Your task is to evaluate the comprehensiveness of this textual description using your chemical molecular knowledge and score it from 1 to 5 according to following standard.

Molecule Description Comprehensiveness Score Standard (1-5)

- Insufficient:**
 - Description provides very minimal or no information about the molecule.
 - Omits key details, making it challenging to form a clear understanding of the molecule.
 - Offers no context or additional properties about the molecule.
- Basic:**
 - Provides a superficial overview, perhaps just mentioning the name or one primary feature.
 - Lacks depth in the description and may not touch on multiple aspects of the molecule.
 - May hint at the molecule's significance but does not provide detailed context.
- Moderate:**
 - Gives a clear description of the molecule's key features, but there might be areas that lack depth.
 - Covers some aspects in detail but may not be exhaustive in its breadth.
 - Provides context about the molecule's properties and significance but may not delve into lesser-known attributes.
- Detailed:**
 - Offers a deep dive into the molecule, detailing its primary features and more.
 - Contains minute omissions but provides a well-rounded understanding of the molecule.
 - Explores various aspects of the molecule, from its structure to its significance in specific fields.
 - Presents a broader context, possibly mentioning its interactions, relevance in research, or other pertinent facts.
- Highly Comprehensiveness:**
 - Delivers a holistic understanding, leaving no stone unturned.
 - Describes not only the key features but also dives its nuances, history, relevance, and lesser-known attributes.
 - Makes use of standard nomenclature, and possibly includes alternative naming or related molecules.
 - Paints a full picture of the molecule's significance, interactions, uses, and any other relevant details.

Input: SMILES: [SMILES]
Description: [Molecular text]

Output:

Figure 5: The prompt template for textual description enrichment and quality evaluation.

pretraining. Note that given the relatively small scale of the Q-former and the associated Uni-Mol model, stage 1 does not involve LoRA tuning.

Stage 2 Pretrain. Initiating from the Stage-1 checkpoint, we proceed to the pretraining of Stage 2 for a text generation training of 10 epochs, maintaining the optimizer and scheduler configurations from the prior stage. To reduce CUDA memory usage, we integrate LoRA with parameters set to $r = 8, \alpha = 32, \text{dropout} = 0.1$. This integration is applied to the modules of [k_proj, v_proj, q_proj, o_proj, gate_proj, up_proj, down_proj]. Utilizing this configuration, the resultant LoRA adapter comprises 19M trainable parameters, representing a mere 0.29% of the total parameter count in Llama2-7B. With a batch size of 16, maximal token length of 320, and accumulated gradients from 2 batches, the computation overhead is 60 GPU hours on NVIDIA A100 with BFloat16 Mixed precision.

Stage 3: Generalist Model. To obtain one checkpoint for all tasks encompassing molecule captioning, descriptive, non-3D-centric property, and 3D-centric property open-text QA, we mix the training

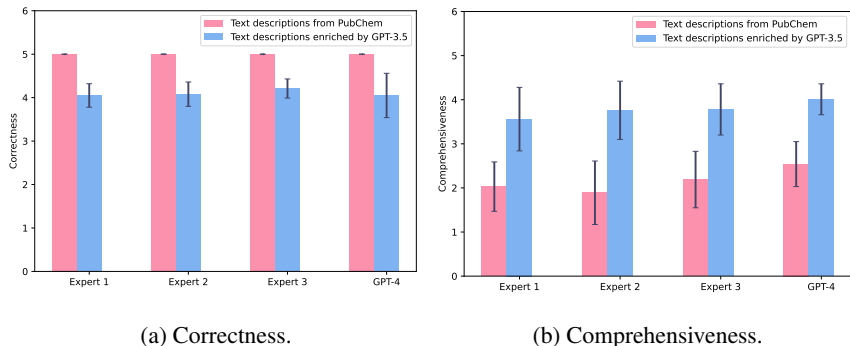


Figure 6: Correctness and comprehensiveness evaluation for subsets of original PubChem texts and GPT-3.5 enriched ones. Figure 5 shows score standards.

sets of these tasks. Given the substantial difference in the size of each dataset, we adopt sampling datasets with probabilities that are proportional to the fourth root of their sizes as in (Dai et al., 2023). Subsequently, the model undergoes training for 400,000 steps with a validation interval of 50,000 steps, utilizing identical optimizer, learning rate scheduler, and LoRA configurations as in stage 2 pretraining. Following this, the checkpoint with the lowest validation loss is evaluated on the test sets of all tasks. The performance of this model is presented in Section 3 as a generalist.

Stage 3: Specialist for Molecule Captioning. The Stage-2 checkpoint is finetuned for 10 epochs and evaluated on the downstream subsets (*i.e.*, train/validation/test) of PubChem with the same optimizer, learning rate scheduler, and LoRA configurations as stage 2 pretraining, with the exception of the warmup steps, which are set to 200,.

Stage 3: Specialist for Open-text QA. We possess three specialists for descriptive, non-3D-centric property, and 3D-centric property Open-text-QA, respectively. Initially, they universally undergo a comprehensive epoch of all instruction tuning data in 3D-MoIT, employing the same configurations as stage 2, but limited to a single epoch. The computational overhead amounts to 200 GPU hours on NVIDIA A100, utilizing BFloat16 Mixed precision. Subsequently, the specialists for descriptive, non-3D-centric property and 3D-centric property Open-text-QA are fine-tuned on the sub-dataset of the corresponding task for 5, 5, and 1 epoch, respectively.

D FAILURE CASE STUDY

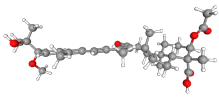
Table 6a shows the mistakes 3D-MoLM made in the molecule captioning task, and we offer the subsequent analysis:

- 3D-MoLM incorrectly identifies Globostellatic acid C as Globostellatic acid B, which are structurally distinct solely in the placement of a methoxy branch. As illustrated in Table 6(b), they have remarkably similar records on PubChem. This suggests that despite accurately identifying primary functions and structures, the discernment of 3D-MoLM at a more refined granularity requires enhancement. We will persist in investigating this in future studies.
- 3D-MoLM inaccurately enumerates the number of atoms in Globostellatic acid C as $C_{34}H_{48}O_6$, whereas the actual composition is $C_{33}H_{48}O_7$. Although this may seem trivial to even non-expert humans, counting atom numbers and providing chemical formulas pose a common challenge for language models. As demonstrated in Figure 7, even the powerful GPT-4 model lacks the ability to directly interpret SMILES strings and provide correct chemical formulas.

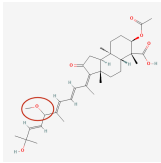
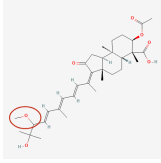
To rectify these deficiencies, we will persist in refining 3D-MoLM from various perspectives, including 1) integrating a more powerful 3D molecular encoder to distinguish molecules with similar structures, 2) investigating more effective projector structures to bridge the gap between 3D molecular representation and LM space, and 3) curating more fine-grained datasets for molecule-text alignment and molecule-centric instruction tuning.

Table 6: Failure case study.

(a) Case of molecule captioning. Mistakes are highlighted.

Molecule	Ground Truth	3D-MoLM
	Globostellatic acid C is a tricyclic triterpenoid of the isomalabaricane group. It has a role as an antineoplastic agent and a metabolite. It is a tricyclic triterpenoid, an acetate ester, an enone, an ether and an oxo monocarboxylic acid. It is a conjugate acid of a globostellatate C(1-).	Globostellatic acid B is a tricyclic triterpenoid with formula C34H48O6 , originally isolated from Globostellata. It has a role as a metabolite. It is a tricyclic triterpenoid, a monocarboxylic acid and an acetate ester. Globostellatic acid B is a natural product found in Globostellata with data available.

(b) PubChem records of Globostellatic acid C and Globostellatic acid B. Their differences are circled in the 2D molecular graphs below. Here we visualize 2D molecular graphs instead of 3D structures for visual clarity.

Molecule	Synonym	Molecular Formula	Textual description on PubChem
	Globostellatic acid B	C ₃₃ H ₄₈ O ₇	Globostellatic acid B is a tricyclic triterpenoid of the isomalabaricane group. It has a role as an antineoplastic agent and a metabolite. It is a tricyclic triterpenoid, an enone, an oxo monocarboxylic acid, an acetate ester and an ether. It is a conjugate acid of a globostellatate B(1-).
	Globostellatic acid C	C ₃₃ H ₄₈ O ₇	Globostellatic acid C is a tricyclic triterpenoid of the isomalabaricane group. It has a role as an antineoplastic agent and a metabolite. It is a tricyclic triterpenoid, an acetate ester, an enone, an ether and an oxo monocarboxylic acid. It is a conjugate acid of a globostellatate C(1-).

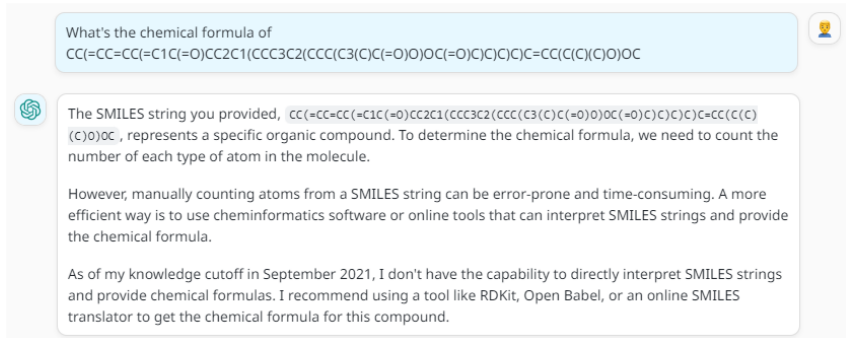


Figure 7: Existing LLMs (e.g., GPT-4) are incapable of accurately counting numbers of atoms and generating chemical formulas of a given SMILES. We use Globostellatic acid B in this example.

E ABLATION STUDY ON INPUT PROMPT

To demonstrate that integrating 3D molecular tokens with 1D SMILES in the prompt enhances molecular understanding, we perform an ablation study and report the results in Table 7. We observe that using a combination of both 3D molecular tokens and 1D SMILES as prompts consistently yields superior performance compared to using only one type. This can be attributed to that including SMILES ensures the LM taps into its inherent understanding of 1D SMILES. Thus, when integrating the 3D molecular encoder and molecule-text projector, the distilled 3D molecular tokens can concentrate on discerning 3D-dependent properties. This is akin to enhancing a model with additional features while retaining its existing representation ability, ensuring a more comprehensive understanding.

Table 7: Ablation study. The instruction prompt of 3D-MoLM contains both 3D molecular tokens and 1D SMILES (*cf.* Figure 1). Llama2-7B, going through stage 3 instruction tuning, is prompted with 1D SMILES, while ‡ denotes prompting with only 3D molecular tokens.

(a) Molecule captioning on the PubChem Dataset.								
Model	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	METEOR		
Llama2-7B	27.01	20.94	35.76	20.68	29.88	32.11		
3D-MoLM‡	27.12	21.10	35.63	20.47	28.84	32.41		
3D-MoLM	29.25	22.07	36.48	21.80	30.95	33.12		

(b) Descriptive Property QA results on 3D-MoLM.								
Model	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	METEOR		
Llama2-7B	27.68	22.81	34.73	21.55	29.91	46.39		
3D-MoLM‡	28.05	23.38	35.47	21.98	30.29	47.20		
3D-MoLM	31.61	25.88	39.93	25.68	34.78	51.74		

(c) Computed property QA results on 3D-MoLM.								
Model	Weight (g/mol)	LogP	TPSA (\AA^2)	Complexity	HOMO (eV)	LUMO (eV)	H-L Gap (eV)	SCF (10^5 eV)
Llama2-7B	27.42 (92%)	1.78 (93%)	17.07 (90%)	78.16 (92%)	1.89 (90%)	1.26 (90%)	1.25 (91%)	0.87 (99%)
3D-MoLM‡	24.15 (93%)	1.42 (93%)	15.02 (92%)	58.84 (91%)	0.68 (94%)	0.52 (93%)	0.48 (92%)	0.51 (99%)
3D-MoLM	16.58 (92%)	0.78 (95%)	10.90 (90%)	45.49 (89%)	0.35 (95%)	0.36 (93%)	0.32 (90%)	0.38 (98%)

Table 8: Ablation study of 3D perception on MolT5.

Base LM	3D Perception	BLEU-2	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	METEOR
MolT5-small	✗	22.53	15.23	30.44	13.45	20.30	23.98
	✓	23.73	16.57	31.57	14.34	21.63	24.81
MolT5-base	✗	24.51	16.61	32.19	14.04	21.35	26.10
	✓	25.55	17.31	33.36	15.49	22.62	27.54
MolT5-large	✗	25.87	17.28	34.07	16.42	23.41	28.04
	✓	26.91	18.50	35.25	17.37	25.50	29.45

F THE EFFECTIVENESS OF 3D PERCEPTION

To demonstrate the effectiveness of 3D perception and eliminate the advantage of model scale (*i.e.*, Llama2 7B versus T5 780M), we integrate the 3D molecule encoder and the projector into MolT5. This facilitates a fair comparison, maintaining the same model size.

As shown in Table 8, MolT5s equipped with 3D perception consistently surpass those lacking 3D perception across all scales. This result substantiates our claim that the incorporation of 3D geometry perception enhances the molecular understanding of language models.

G LIMITATIONS AND FUTURE DIRECTIONS

In this paper, we delve into the domain of 3D molecule-text interpretation, aiming to bridge the understanding between 3D molecular representations and textual descriptions. Although the proposed approach demonstrates promising results on various tasks, we recognize the following limitations:

Expanding to More Molecule-Text Modeling Tasks. Looking forward, tasks like experimental procedure prediction (Vaucher et al., 2021), retrosynthesis (Zhong et al., 2022), and drug-drug interaction (Fang et al., 2024b) can potentially benefit from 3D molecular-text modeling, but are beyond the scope of this work. This limitation stems partly from the lack of openly available datasets for certain tasks and our aim to maintain a focused scope in this study. Additionally, existing tasks like molecular captioning and molecule-text retrieval can be further developed by introducing new datasets that focus on more complex and novel chemical properties (*e.g.*, molecular spectra). Using texts for graph explanation (Li et al., 2022a; Fang et al., 2023a; 2024a) and property prediction (Liu et al., 2023c) can also be explored. These opportunities for expansion are left for future research.

Fine-grained 3D Molecule-Text Alignment. We include case studies in Appendix D showing that 3D-MoLM is limited in discerning fine-grained small molecular structures. Overcoming this issue is

crucial for enhancing performance in tasks that demand a precise understanding of 3D molecular structures. One potential direction is to curate a new 3D molecule-text dataset with explicit 3D coordinate references within textual descriptions. Another direction is to explore more powerful 3D molecular encoders and 3D molecule-text projectors.

Dataset Scale. Other fields of multi-modal modeling have access to datasets comprising more than 10 million paired samples (Changpinyo et al., 2021). In contrast, the datasets in this work (*i.e.*, PubChem 300K and PubChemQC 3M), are relatively limited in scale. This limitation of dataset scales can impede models from reaching its full potential, and underscores the need for larger and better datasets, especially those with high-quality textual descriptions closely related to 3D molecular structures.

Unexplored Capabilities of LMs. Large LMs often exhibit intriguing abilities like in-context learning and chain-of-thought reasoning. However, this research does not venture into leveraging these capabilities, primarily due to the lack of specialized datasets detailing step-by-step inferenced.

Potential Data Leakage. The evaluation of molecular LMs presents another challenge, particularly with large LMs. The vast corpora on which these models are pretrained might inadvertently encompass the test datasets. This overlap can introduce biases, thereby inflating performance metrics.

Future work might address these constraints, ensuring a more holistic and accurate understanding of 3D molecule-text interpretation.