Property Enhanced Instruction Tuning for Multi-task Molecule Generation with Large Language Models

Xuan Lin[1], **Long Chen**[1], **Yile Wang**[2*], **Xiangxiang Zeng**[3], **Philip S. Yu**[4]
[1]School of Computer Science, Xiangtan University
[2]College of Computer Science and Software Engineering, Shenzhen University
[3]College of Information Science and Engineering, Hunan University
[4]Department of Computer Science, University of Illinois
jack_lin@xtu.edu.cn wangyile@szu.edu.cn psyu@cs.uic.edu

Abstract

Large language models (LLMs) are widely applied in various natural language processing tasks such as question answering and machine translation. However, due to the lack of labeled data and the difficulty of manual annotation for biochemical properties, the performance for molecule generation tasks is still limited, especially for tasks involving multi-properties constraints. In this work, we present a two-step framework PEIT (Property Enhanced Instruction Tuning) to improve LLMs for molecular-related tasks. In the first step, we use textual descriptions, SMILES, and biochemical properties as multimodal inputs to pre-train a model called PEIT-GEN, by aligning multi-modal representations to synthesize instruction data. In the second step, we fine-tune existing open-source LLMs with the synthesized data, the resulting PEIT-LLM can handle molecule captioning, text-based molecule generation, molecular property prediction, and our newly proposed multi-constraint molecule generation tasks. Experimental results show that our pre-trained PEIT-GEN outperforms MolT5 and BioT5 in molecule captioning, demonstrating modalities align well between textual descriptions, structures, and biochemical properties. Furthermore, PEIT-LLM shows promising improvements in multi-task molecule generation, proving the scalability of the PEIT framework for various molecular tasks. We release the code, constructed instruction data, and model checkpoints in https://github.com/chenlong164/PEIT.

1 Introduction

Large language models (LLMs) such as GPT-4 (OpenAI, 2023), PaLM (Chowdhery et al., 2023) and LLaMa (Touvron et al., 2023; Dubey et al., 2024) have revolutionized the landscape of artificial intelligence and natural language processing (NLP), allowing machines to understand and generate human language with remarkable fluency and coherence. Based on encoded world knowledge (Petroni et al., 2019) and powerful instruct-following (Zhang et al., 2023) capabilities of LLMs, recent work has successfully used LLM for molecular-related tasks, achieving promising results (Fang et al., 2023; Zhang et al., 2024).

Despite the success, LLMs still have limitations in tasks involving the generation of molecules with restricted properties, therefore limiting its potential applications such as drug discovery (Zhavoronkov, 2018; Elton et al., 2019). The challenges for tackling such tasks mainly lie in three aspects: (1) Existing studies have shown limitations of LLMs in understanding molecular representations (Grisoni, 2023), which makes it more challenging for handling such tasks with precise properties; (2) While there is some known SMILES-property pairing data, it often remains limited to predicting a single property and lacks datasets encompassing a wide range of properties (Wu et al., 2018). Moreover, most of these datasets do not include precisely described textual data, making it challenging to identify accurate tri-modal data pairs (Krenn et al., 2020); (3) To our knowledge, there are no suitable datasets or evaluation methods for multi-constraint molecule generation using LLMs, which poses challenges in standardizing and assessing such molecule generation tasks with these models (Jin et al., 2018; Elton et al., 2019).

To address these challenges, we propose a framework called PEIT (Property Enhanced Instruction Tuning) to generate multi-modal molecular instruction datasets in bulk, aiming to enhance the capabilities of LLMs in multi-task molecule generation. Using the PEIT framework, our pre-trained model can handle both general tasks (e.g., molecule captioning (Edwards et al., 2022)) and property-related tasks such as property prediction (Chang and Ye, 2024). This makes it suitable for constructing data to evaluate multi-constraint molecule generation

*Corresponding author.

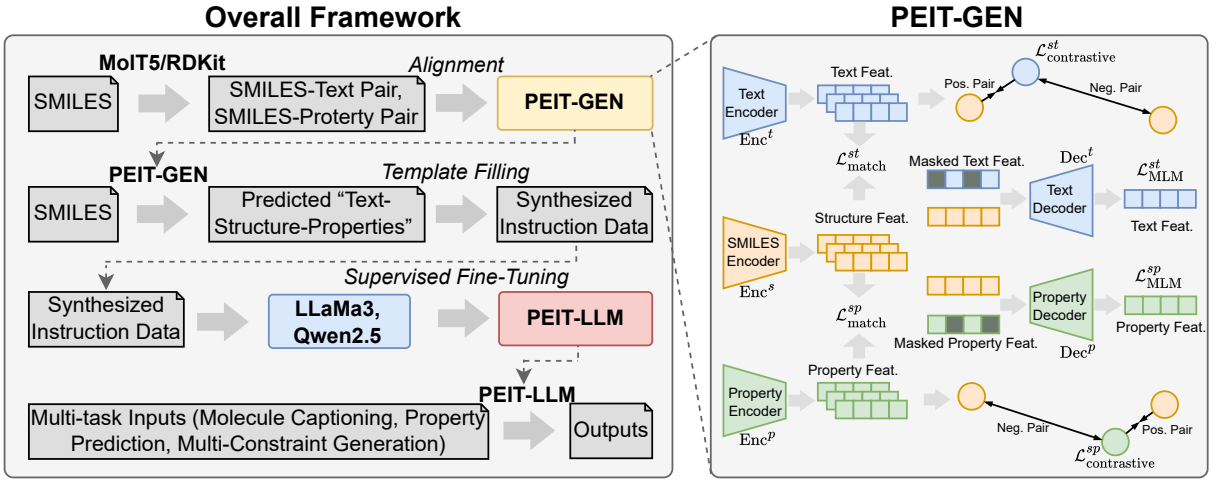

Figure 1: Left: Overall PEIT framework. We first pre-train the PEIT-GEN and construct instruction data via template filling. Then we fine-tune the open-source LLMs through instruction tuning, the resulting PEIT-LLM is used for multi-task molecule generation. Right: The process of PEIT-GEN pre-training, see details in section 3.2.

capabilities and for serving as instruction tuning data to improve existing open-source LLMs.

The overall structure of the proposed PEIT framework is shown in the left of Figure 1. Specifically, it consists of two components: (1) We pre-train a model called PEIT-GEN through multi-modal representation alignment, which integrates text-based (molecular descriptions), structure-based (SMILES), and property-based (property-value pairs) information to generate diverse unstructured text, sequence, and property data; (2) By using the synthesized instruction data, we fine-tune open-source LLMs and develop PEIT-LLM, which can be applied to various molecule generation tasks mentioned above, including our proposed multi-constraint molecule generation, which simulates real-world drug discovery scenarios where the generation of drug molecules is guided by multiple properties constraints.

Experimental results demonstrate that our pre-trained PEIT-GEN achieves competitive or better results in molecule captioning tasks, comparing to a variety of biomolecular models including MolT5 (Edwards et al., 2022), BioT5 (Pei et al., 2023), GIT-Mol (Liu et al., 2024), MolXPT (Liu et al., 2023b). Additionally, PEIT-LLM based on LLaMa3.1-8B (Dubey et al., 2024) exhibits superior performance compared to both specialized models (BioT5+ (Pei et al., 2024) and Mol-Instructions (Fang et al., 2023)) and general-purpose LLMs (LLaMa3 (Touvron et al., 2023), LLaMa3.1 (Dubey et al., 2024), and Qwen2.5 (Yang et al., 2024)) in molecular property prediction and our newly proposed multi-constraint molecule generation tasks.

Our contributions can be summarized as follows:

- We propose PEIT, a novel framework that enables existing open-source LLMs to align the textual descriptions, SMILES sequences, and biochemical properties through multi-modal representation alignment, thereby facilitating multi-task molecule generation.
- PEIT achieves promising results in various benchmarks. It surpasses the baselines at least by 2.3% on BLEU-2 for molecule captioning task. Moreover, in text-based molecule generation task, PEIT shows a considerable advantage of 21.76 Levenshtein over the baselines. For five-property constraint molecule generation, PEIT outperforms the baselines in terms of RMSE and R^2, respectively.
- We conduct ablation studies to show PEIT's effectiveness of incorporating multi-modal features into existing open-source LLMs for multi-task molecule generation. Additionally, our quantitative analysis shows that various modeling choices, including different objectives and SFT steps, help improve the LLM's potential to accelerate the drug discovery.

2 Related Work

Molecule generation. The goal is to design and generate molecules that meet specific properties, deep learning models have emerged and they are mainly categorized as follows: (1) text-based

molecule generation that uses textual descriptions to generate molecules that match the given description (Liu et al., 2023b, 2024). MolT5 (Edwards et al., 2022) was the first proposed to realize translation between textual description and molecular SMILES. BioT5 aims to enhance molecular understanding by incorporating protein modality. They also perform molecule captioning, which is equivalent to the inverse task of text-based molecule generation. (2) property-guided molecule generation is the inverse process of molecular property prediction, where molecules are generated based on specific biochemical property constraints. Notably, SPMM (Chang and Ye, 2024) was the first to establish a connection between 53 biochemical properties and SMILES sequences, making multi-constraint molecule generation possible. However, few existing models can simultaneously perform text-based or multi-constraint molecule generation and molecule captioning.

Molecular property prediction. Deep learning models have been developed for molecular property prediction each with their own advantages and limitations. Transformer-based models design attention mechanism to capture contextual contexts from large-scale SMILES sequences (Ross et al., 2022). The molecular graph can be directly obtained from SMILES sequences via RDKit (Landrum et al., 2013). Graph-based models develop diverse graph neural networks to learn differentiable representations (Wang et al., 2022). However, these methods ignore the potential that incorporating textual knowledge enables to realize new drug design objectives (Zeng et al., 2022; Liu et al., 2023a). Recently, a novel molecular pre-trained model named SPMM (Chang and Ye, 2024) that extends the application of multimodal pre-training approaches by aligning molecular structures and biochemical properties. This paper extends the multimodal pre-training to patterns of text-sequence-property triplets, which is defined flexibly by LLM-understandable textual prompts.

Instruction tuning. Specialized datasets construction seems the effective way to enable LLMs to better perform the molecular-related tasks. For instance, Mol-Instructions (Fang et al., 2023) provides a large-scale biomolecular instruction dataset designed for LLMs, which contains a variety of instruction data ranging from small molecules, proteins, and biomolecular texts. Additionally, BioT5+ (Pei et al., 2024) integrates IUPAC names, extensive biological texts, and molecular data through multi-task instruction tuning, providing more comprehensive insights in the fields of drug discovery. How to generate reliable data related to molecular knowledge remains a challenge of instruction tuning for existing open source LLMs.

3 Method

3.1 Overview of PEIT Framework

The overview of PEIT framework is shown in Figure 1 (left), which consists of PEIT-GEN and PEIT-LLM. In PEIT-GEN, we generate a large number of "SMILES-text" and "SMILES-property" pairs to serve as multi-modal data. Then we design multiple multi-modal alignment objectives to pre-train PEIT-GEN. In PEIT-LLM, by using the pre-trained PEIT-GEN, we can predict a large number of triplets to generate more diverse SMILES inputs, and then construct diverse instruction data based on template filling. By utilizing the synthesized instruction data, PEIT-LLM enables the supervised fine-tuning of open-source LLMs including LLaMa (Dubey et al., 2024) and Qwen (Yang et al., 2024), enhancing the capabilities for multi-task molecule generation.

3.2 Pre-training of PEIT-GEN

The pre-training stage of PEIT-GEN is shown in the right of Figure 1. For a given molecule, different representations offer unique and complementary features, which are crucial for comprehensive molecule understanding. PEIT-GEN aims to integrate information from three modalities simultaneously, including textual information $\mathcal{T}$ (text), molecular structure $\mathcal{S}$ (SMILES), and biochemical properties $\mathcal{P}$ (property-value). Such ability can help synthesizing sufficient instruction data for further enhancing the ability of LLMs. In particular, PEIT-GEN consists of three Transformer encoders Enc^t, Enc^s, Enc^p and two decoders Dec^t, Dec^p, and we design different training objectives to align features from different modalities.

Cross-modal representation matching. We leverage pre-trained models SciBERT (Beltagy et al., 2019) as trainable Enc^t for encoding textual data, BERT (Devlin et al., 2019) as Enc^s and Enc^p for encoding SMILES and properties. Then we obtain feature representations across all three modalities, establishing the foundation for feature alignment.

We propose cross-modal representation matching to align the representations from different perspectives by the same molecule. In particular, we introduce the SMILES-text matching loss $\mathcal{L}^{st}_{\text{match}}$

and the SMILES-property matching loss $\mathcal{L}^{sp}_{\text{match}}$, which serve as objectives for training the encoders. In this way, the model can effectively learn cross-modal relationships and improve performance in multi-modal tasks by aligning the feature spaces. The matching loss is calculated as follows:

$$\mathcal{L}^{st}_{\text{match}} = \ell_{\text{CE}}\left(y^{st}_{\text{match}}, \text{MLP}(\text{Enc}^s(\mathcal{S}) \oplus \text{Enc}^t(\mathcal{T}))\right), \tag{1}$$

$$\mathcal{L}^{sp}_{\text{match}} = \ell_{\text{CE}}\left(y^{sp}_{\text{match}}, \text{MLP}(\text{Enc}^s(\mathcal{S}) \oplus \text{Enc}^p(\mathcal{P}))\right), \tag{2}$$

where y^{st}_{match} and y^{sp}_{match} are labels as 0 or 1, indicating whether the corresponding SMILES-text or SMILES-property pairs are matching. $\text{Enc}(\cdot)$ indicates the representation of the data (i.e., [CLS] token of Transformer encoder), $\oplus$ is the concatenation operation, and $\text{MLP}(\cdot)$ is the trainable multi-layer perception. The encoders are optimized by the cross-entropy loss ℓ_{CE} using the given data from different modalities.

Multi-modal contrastive learning. The representation matching can be viewed as an explicit 2-way classification training. We further utilize contrastive learning to directly enhancing the representation by pulling semantically close neighbors together and pushing apart non-neighbors from data of different modalities. To calculate the similarity between the encoded features of different modalities, we extract the encoded features and then compute the instance-level similarities through the inner product:

$$\text{sim}(\mathcal{S}, \mathcal{T}) = \left(\text{MLP}^s(\text{Enc}^s(\mathcal{S}))\right)^\top \text{MLP}^t(\text{Enc}^t(\mathcal{T})), \tag{3}$$

$$\text{sim}(\mathcal{S}, \mathcal{P}) = \left(\text{MLP}^s(\text{Enc}^s(\mathcal{S}))\right)^\top \text{MLP}^p(\text{Enc}^p(\mathcal{P})), \tag{4}$$

where MLP^s, MLP^t and MLP^p are multi-layer perceptions applied to SMILES, text, and property representations, respectively. Then, for the given SMILES $\mathcal{S}$, text $\mathcal{T}$, and property $\mathcal{P}$, we compute the cross-modal batch-level similarities as follows:

$$s_{s2t} = \frac{\exp(\text{sim}(\mathcal{S},\mathcal{T})/\tau)}{\sum_{i=1}^{M} \exp(\text{sim}(\mathcal{S},\mathcal{T}_i)/\tau)}, \tag{5}$$

$$s_{s2p} = \frac{\exp(\text{sim}(\mathcal{S},\mathcal{P})/\tau)}{\sum_{i=1}^{N} \exp(\text{sim}(\mathcal{S},\mathcal{P}_i)/\tau)}, \tag{6}$$

where M and N represent the total number of texts and property in the batch of data pairs, respectively. τ is the temperature controlling the sharpness of the similarity. The intra-modal similarities s_{s2s}, s_{p2p}, and s_{t2t} can be computed in similar manners.

Based on the cross-modal and intra-modal batch-level similarities, the contrastive loss is formulated by calculating the cross-entropy according to one-hot encoded similarity vectors y, where the value is 1 for pairs derived from the same molecule or 0 for all other combinations:

$$\begin{aligned}\mathcal{L}^{st}_{\text{contrastive}} = \frac{1}{2}(&\ell_{\text{CE}}(y_{s2t}, s_{s2t}) + \ell_{\text{CE}}(y_{t2s}, s_{t2s}) \\ &+ \ell_{\text{CE}}(y_{s2s}, s_{s2s}) + \ell_{\text{CE}}(y_{t2t}, s_{t2t})),\end{aligned} \tag{7}$$

$$\begin{aligned}\mathcal{L}^{sp}_{\text{contrastive}} = \frac{1}{2}(&\ell_{\text{CE}}(y_{s2p}, s_{s2p}) + \ell_{\text{CE}}(y_{p2s}, s_{p2s}) \\ &+ \ell_{\text{CE}}(y_{s2s}, s_{s2s}) + \ell_{\text{CE}}(y_{p2p}, s_{p2p})).\end{aligned} \tag{8}$$

Cross-modal masked language modeling. To further strengthen the model's capability in molecule captioning, we employ the masked language modeling (MLM; Devlin et al., 2019) to enhance the model performance on text generation. MLM is originally designed for the BERT encoder, which is not specifically used for generation. We design decoders to generate original unmasked property and textual description sequences, under the guidance of SMILES features through cross-attention. Specifically, given a pair of text and property, the calculation of vanilla self-attentions are as follows:

$$\begin{aligned}\text{SelfAtt}(\mathcal{T}) &\doteq \text{softmax}(W^t_Q h(\mathcal{T})(W^t_K h(\mathcal{T}))^\top) W^t_V h(\mathcal{T}), \\ \text{SelfAtt}(\mathcal{P}) &\doteq \text{softmax}(W^p_Q h(\mathcal{P})(W^p_K h(\mathcal{P}))^\top) W^p_V h(\mathcal{P}),\end{aligned} \tag{9}$$

where $h(\cdot)$ denotes the hidden representations, W_Q, W_K, and W_V are the matrix for query, key, and values among the same modality, respectively.

For text decoder Dec^t and property decoder Dec^p, we propose cross-modal MLM objectives which further integrates SMILES features for masked text or property prediction via applying cross-attention:

$$\begin{aligned}\text{CrossAtt}(\mathcal{T}) &\doteq \text{softmax}(W^t_Q h(\mathcal{T})(W^s_K h(\mathcal{S}))^\top) W^t_V h(\mathcal{T}), \\ \text{CrossAtt}(\mathcal{P}) &\doteq \text{softmax}(W^p_Q h(\mathcal{P})(W^s_K h(\mathcal{S}))^\top) W^p_V h(\mathcal{P}).\end{aligned} \tag{10}$$

By introducing the SMILES features in attention layers for MLM training, we enable the model to utilize SMILES-text and SMILES-property data pairs to perform molecule captioning and property prediction. The cross-modal MLM loss $\mathcal{L}^{st}_{\text{MLM}}$ and $\mathcal{L}^{sp}_{\text{MLM}}$ are computed as follows:

$$\mathcal{L}^{st}_{\text{MLM}} = -\sum_{i=1}^{N}\sum_{j=1}^{n} \log \text{Prob}\left(w^{(i)}_j \mid \text{Dec}^t(\tilde{\mathbf{w}}^{(i)}_{\neg j}); \theta_t\right), \tag{11}$$

$$\mathcal{L}^{sp}_{\text{MLM}} = -\sum_{i=1}^{N}\sum_{j=1}^{n} \log \text{Prob}\left(w^{(i)}_j \mid \text{Dec}^p(\tilde{\mathbf{w}}^{(i)}_{\neg j}); \theta_p\right), \tag{12}$$

Task	MolT5	BioT5	BioT5+	MolXPT	Git-Mol	SPMM	LLaMa, Qwen	PEIT-LLM (ours)
Molecule Captioning	✓	✓	✓	✓	✓	✗	✓ (limited)	✓
Text-Based Molecule Generation	✓	✓	✓	✓	✓	✗	✓ (poor)	✓
Molecular Property Prediction	✗	✗	✗	✗	✓	✓	✓ (poor)	✓
Multi-Constraint Molecule Generation	✗	✗	✗	✗	✗	✗	✓ (poor)	✓

Table 1: Comparing PEIT-LLM with biomolecular models and general-purpose LLMs on molecular-related tasks.

where Prob is the conditional probability to predict the word $w_j^{(i)}$ in the vocabulary, N is the total number of samples, n is the number of masked words in each sample, $\tilde{\mathbf{w}}_{\neg j}^{(i)}$ is the sequence after masking the j-th word in the i-th sample, θ_t and θ_p are the trainable parameters corresponding to each modality in the two decoders. For the masked language training, we adapt a step-by-step strategy on generation task, where each step predicts the next token based on the generated contexts. In each iteration, the model calculates the conditional probability of each candidate token and then selects the optimal one as the output. In this way, our model can achieve complete cross-modal generation through multiple iterations.

Training. The overall training objective for pre-training PEIT-GEN is to minimize the sum of all three types of losses across three modalities:

$$\mathcal{L} = \mathcal{L}^{st}_{\text{match}} + \mathcal{L}^{sp}_{\text{match}} + \alpha\mathcal{L}^{st}_{\text{contrastive}} + \alpha\mathcal{L}^{sp}_{\text{contrastive}} + \beta\mathcal{L}^{st}_{\text{MLM}} + \beta\mathcal{L}^{sp}_{\text{MLM}}, \quad (13)$$

where α and β are hyper-parameters for balancing different loss terms.

3.3 Instruction Tuning for PEIT-LLM

Template Filling. The pre-trained PEIT-GEN offers unstructured data in the format of "text-SMILES-properties" (i.e., text-structure-property) triplets. To obtain more task-specific data and to adapt to the strong instruction-following abilities of LLMs, we design templates for different downstream tasks, as shown in Figure 5 in Appendix A. For instance, in the text-based molecule generation task, we fix a general question format and then extract molecular descriptions from unstructured data to fill the description part of a pre-defined template, resulting in a natural question as instructions. The SMILES from unstructured triplets is used as the desired response. In this way, we can generate diverse task-specific instruction data in bulk for subsequent instruction tuning for LLMs.

Multi-constraint molecule generation task. Molecule generation often requires to be conducted under multiple constraints rather than a single condition. In this work, we propose a new task to assess molecule generation through a variety of descriptors, by comparing the alignment between the generated molecules and specific criteria to evaluate the generative performance of LLMs. By using the large-scale unstructured data generated by PEIT-GEN, we can effectively synthesize sufficient data for evaluation. Specifically, we follow SPMM (Chang and Ye, 2024) and predict 5 common properties out of the 53 available biochemical properties for diverse SMILES, including ExactMolWt, MolLogP, MolWt, QED, and TPSA. Based on the template filling, the predicted multiple property-values can be used to construct data for multi-constraint molecule generation.

Supervised fine-tuning. We select LLaMa3.1-8B (Dubey et al., 2024) and Qwen2.5-7B (Yang et al., 2024) as base LLMs. We then perform standard supervised fine-tuning (SFT; Ouyang et al., 2024) by using the "instruction-response" pairs. In practice, we construct totally 1 million instruction data of four different tasks (i.e., molecule captioning, text-based molecule generation, property prediction, and multi-constraint molecule generation) from 200k unstructured "text-SMILES-properties" triplets obtained by PEIT-GEN.

3.4 Comparing PEIT-LLM with Biomolecular Models and LLMs

Table 1 shows a comparison of our PEIT-LLM with existing pre-trained models and general LLMs on multiple molecular generation tasks. For most of the pre-trained models such as MolT5 and BioT5, they focus on molecule captioning and text-based molecule generation, which can not handle property-related tasks. SPMM is a specialized model for property prediction. However, it lacks of generation ability due to the lack of textual descriptions. Current LLMs such as LLaMa and Qwen show strong performance on general NLP-based tasks through conversations or instruction-following. However, these general LLMs still have limitations in tasks related to molecule generation due to a lack of molecular knowledge. In contrast, through fine-tuning on diverse instruction data with

Model	Data Size ↓	BLEU-2 ↑	BLEU-4 ↑	METEOR ↑	ROUGE-1 ↑	ROUGE-2 ↑	ROUGE-L ↑
MolT5-small (Edwards et al., 2022)	100M	0.513	0.398	0.492	0.567	0.412	0.501
MolT5-large (Edwards et al., 2022)	100M	0.594	0.508	0.613	0.654	0.508	0.592
BioT5 (Pei et al., 2023)	33M	**0.635**	**0.556**	0.656	0.692	0.559	0.633
GIT-Mol (Liu et al., 2024)†	4.8M	0.352	0.263	0.533	0.575	0.485	0.560
MolXPT (Liu et al., 2023b)†	30M	0.594	0.505	0.626	0.660	0.511	0.597
PEIT-GEN (ours)	**0.48M**	0.598	0.534	**0.676**	**0.700**	**0.582**	**0.653**

Table 2: Results on CHEBI-20 molecule captioning with different pre-trained models. †Results are reported from papers accordingly. The best results in each column are **in bold**, and the second-best results are underlined.

rich molecular knowledge, PEIT-LLM can perform multiple molecule generation tasks simultaneously.

4 Experiments

4.1 Experimental Setup

Dataset. For pre-training PEIT-GEN, we extract approximately 480k molecular SMILES entries from the ZINC dataset (Irwin et al., 2012) and then generate SMILES-text pair data using MolT5 (Edwards et al., 2022). Additionally, we calculate 53 biochemical property-value via RDKit, resulting in nearly 480k "text-SMILES-properties" triplets for pre-training. Following MolT5, we use the CHEBI-20 dataset (Edwards et al., 2021) to evaluate PEIT-GEN's performance on molecule captioning and molecular property prediction. We split the CHEBI-20 dataset into training, validation, and test sets with an 8:1:1 ratio, and we verify the property values of each molecule via RDKit.

For pre-training PEIT-LLM, we utilize the 200k tri-modal data generated by PEIT-GEN and employ template filling to generate 200k instruction data for each downstream task. For molecular property prediction, we select two biochemical properties with distinct differences for evaluation, generating 200k instruction data for each property. Finally, we obtain a total of 1000k instruction data across four tasks for SFT training. Similar to PEIT-GEN, molecular property prediction tasks on PEIT-LLM can be validated by RDKit on CHEBI-20 dataset.

Baseline Models. To demonstrate the efficacy of PEIT-GEN and PEIT-LLM, we compare various popular pre-trained models and LLMs including MolT5 (Edwards et al., 2022), BioT5 (Pei et al., 2023), BioT5+ (Pei et al., 2024), MolXPT (Liu et al., 2023b), GIT-Mol (Liu et al., 2024), SPMM (Chang and Ye, 2024), LLaMa3 (Touvron et al., 2023), LLaMa3.1 (Dubey et al., 2024), Qwen2.5 (Yang et al., 2024), and Mol-Instructions (Fang et al., 2023). Details of these baselines and evaluation metric are in Appendix B and C, respectively.

Model	Modality	Data Size ↓	R^2 ↑	RMSE ↓
SPMM (Chang and Ye, 2024)	$\mathcal{S}, \mathcal{P}$	1.5M	**0.921**	0.194
PEIT-GEN (ours)	$\mathcal{S}, \mathcal{P}, \mathcal{T}$	480K	0.910	**0.169**

Table 3: Comparing performance of our PEIT-GEN to SPMM on molecular property prediction.

Implementation Details. For pre-training PEIT-GEN, we follow BERT-base (Devlin et al., 2019) and set hidden size as 768 for three encoders. The training batch is 16, temperature τ is 0.07, the mask probability for computing $\mathcal{L}_{\text{MLM}}$ is 0.15, and the momentum parameter is 0.995 with AdamW optimizer (Loshchilov, 2017). We pre-train PEIT-GEN with 20 epochs and then fine-tune it on CHEBI-20 training set for 200 epochs, with a learning rate of 5e-4. For supervised fine-tuning PEIT-LLM, we use LLaMa-Factory (Zheng et al., 2024) framework and apply LoRA (Hu et al., 2022) fine-tuning for 6 epochs. The maximum length is 1024, batch size is 3, and the initial learning rate to 5e-5 with cosine decay. For baselines, we use the released checkpoints for evaluation. All experiments are run on NVIDIA 4090 GPUs with 24GB memory.

4.2 Comparing PEIT-GEN with Pre-trained Biomolecular Models

Molecule captioning. Results on CHEBI-20 molecule captioning are shown in Table 2. Our model demonstrates superior performance in generating high-quality and relevant molecular caption. PEIT-GEN achieved the best results in METEOR and ROUGE, and the second-best performance in BLEU. Notably, compared to BioT5 which performs the best in BLEU, our approach requires significantly less data. This indicates that using domain-specific models to generate paired data for pre-training is more efficient than single-modality pre-training, enabling excellent performance with much less training data.

Molecular property prediction. The performance of PEIT-GEN in molecular property prediction is shown in Table 3. Following SPMM, we evaluate

Model	#Params	BLEU-2 ↑	BLEU-4 ↑	METEOR ↑	ROUGE-1 ↑	ROUGE-2 ↑	ROUGE-L ↑
LLaMa3 (Touvron et al., 2023)	7B	0.032	0.002	0.117	0.121	0.010	0.065
LLaMa3.1 (Dubey et al., 2024)	8B	0.042	0.004	0.121	0.140	0.019	0.095
Qwen2.5 (Yang et al., 2024)	7B	0.049	0.007	0.188	0.177	0.029	0.112
BioT5+ (Pei et al., 2024)	5.4B	**0.774**	**0.732**	**0.804**	**0.825**	**0.753**	**0.800**
Mol-Instructions (Fang et al., 2023)	8B	0.217	0.143	0.254	0.337	0.196	0.291
PEIT-LLM-Qwen2.5 (ours)	7B	0.422	0.314	0.468	0.535	0.361	0.477
PEIT-LLM-LLaMa3.1 (ours)	8B	0.425	0.316	0.475	0.541	0.370	0.489

Model	#Params	BLEU ↑	Validity ↑	Levenshtein ↓	MACCS FTS ↑	Morgan FTS ↑	RDKit FTS ↑
LLaMa3 (Touvron et al., 2023)	7B	0.261	0.330	45.788	0.372	0.127	0.213
LLaMa3.1 (Dubey et al., 2024)	8B	0.270	0.368	43.183	0.411	0.138	0.248
Qwen2.5 (Yang et al., 2024)	7B	0.217	0.245	50.550	0.403	0.110	0.276
BioT5+ (Pei et al., 2024)	5.4B	0.701	**1.000**	39.790	0.864	**0.703**	0.764
Mol-Instructions (Fang et al., 2023)	8B	0.345	**1.000**	41.367	0.412	0.147	0.231
PEIT-LLM-Qwen2.5 (ours)	7B	0.810	0.950	21.133	0.832	0.619	0.735
PEIT-LLM-LLaMa3.1 (ours)	8B	**0.836**	0.970	**18.030**	**0.875**	0.661	**0.776**

Table 4: Results on molecule captioning and text-based molecule generation with different LLMs.

on 1,000 molecules from the ZINC dataset which were not included in the training set. Compared to SPMM, which is specifically designed for property prediction, PEIT-GEN achieves comparable performance while using only one-third of the data size across three modalities. In terms of RMSE, PEIT-GEN outperformed SPMM, while SPMM was slightly ahead by 0.11 percentage points in the R^2 metric. These results demonstrate that PEIT-GEN can generate high-quality biochemical properties of molecules, highlighting the critical role of high-quality multimodal data in advancing molecular understanding tasks.

4.3 Comparing PEIT-LLM with LLMs

Molecule captioning.

As shown in Table 4, the comparison results show that our model outperforms general-purpose LLMs (Qwen-2.5 and LLaMa3.1) as well as Mol-Instructions, which utilizes a biochemical information instruction dataset for SFT. PEIT-LLM achieves the second-best performance in BLEU, METEOR, and ROUGE, but still lags behind BioT5+, which is specifically trained for molecule captioning task. This indicates that the responses from BioT5+ are closer to the standard answers of CHEBI-20, while PEIT-LLM generates more diverse responses. By comparing with Mol-Instructions, we demonstrate the quality of generated data by PEIT-GEN and the effectiveness of our instruction data through multi-task template-filling. Case study is provided in Table 6 of Appendix E to further illustrate this point.

Text-based molecule generation. The results for text-based molecule generation on the CHEBI-20 test set are shown in Table 4. PEIT-LLM outperforms other baselines in numerical metrics such as BLEU score, Levenshtein Distance, MACCS Fingerprint Similarity, Morgan Fingerprint Similarity, and RDKit Fingerprint Similarity. Meanwhile, BioT5+ and Mol-Instructions show an advantage in the Validity metric. This indicates that PEIT-LLM, after multi-task instruction fine-tuning, has a strong understanding of the key structural representations of molecules as well as their textual descriptions. Case study is provided in Table 7 of Appendix E to further illustrate this point. This also indirectly validates the effectiveness of the instruction data synthesized by our proposed PEIT-GEN.

Model	MolWt PP	MolLogP PP	Five-Property CG	
	(RMSE) ↓	(RMSE) ↓	(RMSE) ↓	(R^2) ↑
LLaMa3 (Touvron et al., 2023)	491.542	561.523	79.125	-0.639
LLaMa3.1 (Dubey et al., 2024)	544.517	552.521	74.646	-0.652
Qwen2.5 (Yang et al., 2024)	100.161	132.141	75.991	-0.967
Mol-Instructions (Fang et al., 2023)	72.172	1.313	71.991	-0.352
PEIT-LLM-Qwen2.5 (ours)	14.164	0.164	19.750	0.550
PEIT-LLM-LLaMa3.1 (ours)	**13.918**	**0.141**	**14.212**	**0.613**

Table 5: Results on MolWt, MolLogP property prediction (PP), and five-property constraint molecule generation (CG) with different LLMs.

Molecular property prediction. For predicting single property, due to the large number of property, we selected two representative ones for prediction. The property ExactMolWt with relatively large numerical values (usually 100∼1000), and property MolLogP with relatively small numerical values (usually -5∼10) are shown in Table 5. The results show that PEIT-LLM outperforms all other LLMs in predicting specific biochemical properties, demonstrating that PEIT-LLM exhibits strong sensitivity to molecular properties, showing excellent predictive performance for both properties with large numerical values and those with

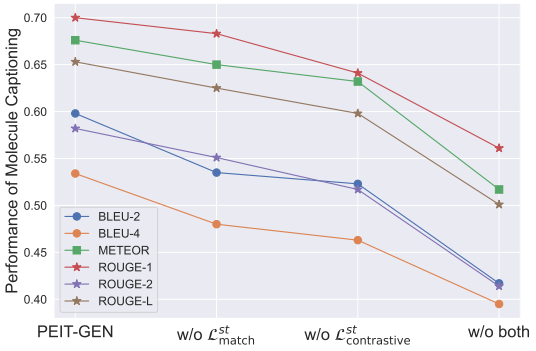

Figure 2: Ablation study of different objectives during PEIT-GEN pre-training.

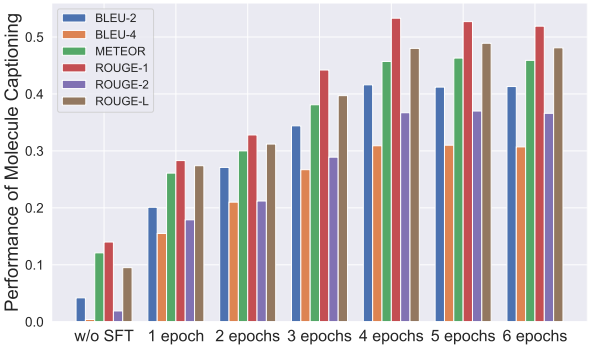

Figure 3: The impact of different amount of SFT steps for PEIT-LLM on molecule captioning.

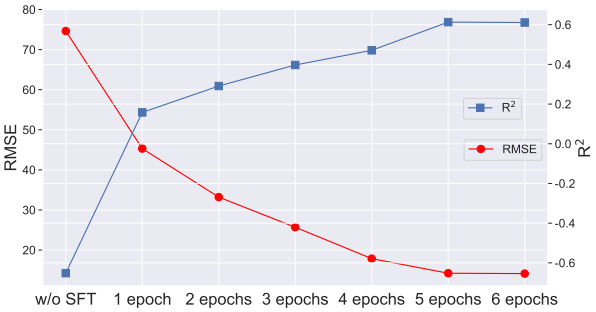

Figure 4: The impact of different amount of SFT steps for PEIT-LLM on multi-constraint molecule generation.

smaller values. This confirms the feasibility of using multi-task SFT to enhance LLMs' understanding of molecular properties and further validates the reliability of the molecular property instruction dataset. Case study is provided in Table 8 of Appendix E to further illustrate this point.

Multi-constraint molecule generation. Results for our proposed multi-constraint molecule generation task is shown in Table 5. PEIT-LLM surpasses baselines by large margin in both RMSE and R^2 metrics. Case study is provided in Table 9 of Appendix E to further illustrate this point. Note that this task requires the model to meet the demands of multiple properties with precise values, placing high demands on the model's overall understanding capability. General-purpose LLMs, or those not specifically trained for this task, lack the required information storage and fitting abilities. As demonstrated, through our property enhanced instruction tuning, the model gain strong molecular understanding capabilities.

4.4 Analysis

Ablation study. Figure 2 shows the ablation study of SMILES-text matching loss $\mathcal{L}^{st}_{\text{match}}$ and cross-modal contrastive loss $\mathcal{L}^{st}_{\text{contrastive}}$, which are not considered in SPMM due to the lack of textual description modality[1]. By removing these training objectives, the performance degradation across all metrics, with a more significant decline when both are removed simultaneously. This demonstrates that both $\mathcal{L}^{st}_{\text{match}}$ and $\mathcal{L}^{st}_{\text{contrastive}}$ are helpful in cross-modal feature alignment, thereby enhancing the performance of molecule captioning.

Impact of SFT steps. Figure 3 and Figure 4 show the results of PEIT-LLM with different SFT steps. We find that the performance steadily improved at first few epochs, showing that the instruction data is useful for both molecule captioning and multi-constraint molecule generation tasks. The performance gradually saturates around epochs 5-6. This indicates that the LLaMa-7B model achieves optimal performance with 1 million instruction data, and further training might lead to over fitting.

[1] $\mathcal{L}^{st}_{\text{MLM}}$ and $\mathcal{L}^{sp}_{\text{MLM}}$ are necessary for caption generation via decoders, thus we do not consider them in ablation study.

5 Conclusion

We propose a novel framework PEIT that aims to enable open-source LLMs to perceive multi-modal features for multi-task molecule generation. For this purpose, PEIT establishes cross-modal connections among molecular structures, textual description, and biochemical properties through multi-modal representation alignment. Through template filling, PEIT can help synthesizing diverse task-specific instruction data for LLMs. We further introduce a new multi-constraint molecule generation task that requires generating novel molecules meeting multiple property constraints. Experiments show that PEIT achieves promising performances on molecule captioning, text-based molecule generation, and property-related tasks compared with various biomolecular models and LLMs.

Limitations

While PEIT is capable of achieving comparative or better performance over existing studies, it still has some limitations as follows: First, PEIT integrates the pre-trained PEIT-GEN model as part of the pipeline, so the performance of PEIT-GEN greatly affect the overall performance of PEIT-LLM. Second, PEIT-GEN uses three types of modality to construct the instruction data. However, some modalities data (e.g., knowledge graph and molecular images) might be more crucial than sequences for the molecular-related task. As a result, exploring the different modalities might lead to a different result. Lastly, the template utilized for instruction-tuning in this work still relies on manual design. Our approach is influenced by previous study that has been shown to be effective. Nevertheless, it would be intriguing to explore the development of automated methods for constructing superior instruction-tuning templates.

References

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.

Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A pretrained language model for scientific text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620.

Adrià Cereto-Massagué, María José Ojeda, Cristina Valls, Miquel Mulero, Santiago Garcia-Vallvé, and Gerard Pujadas. 2015. Molecular fingerprint similarity search in virtual screening. *Methods*, 71:58–63.

Jinho Chang and Jong Chul Ye. 2024. Bidirectional generation of structure and properties through a single molecular foundation model. *Nature Communications*, 15(1):2323.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. 2022. Translation between molecules and natural language. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 375–413.

Carl Edwards, ChengXiang Zhai, and Heng Ji. 2021. Text2mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 595–607.

Daniel C Elton, Zois Boukouvalas, Mark D Fuge, and Peter W Chung. 2019. Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering*, 4(4):828–849.

Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Huajun Chen. 2023. Mol-Instructions: A large-scale biomolecular instruction dataset for large language models. In *International Conference on Learning Representations*.

Francesca Grisoni. 2023. Chemical language models for de novo drug design: Challenges and opportunities. *Current Opinion in Structural Biology*, 79:102527.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

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):1757–1768.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2018. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pages 2323–2332.

Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. 2020. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024.

Greg Landrum et al. 2013. Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum*, 8(31.10):5281.

VI Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Proceedings of the Soviet physics doklady*.

Pengfei Liu, Yiming Ren, Jun Tao, and Zhixiang Ren. 2024. Git-mol: A multi-modal large language model for molecular science with graph, image, and text. *Computers in biology and medicine*, 171:108073.

Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Animashree Anandkumar. 2023a. Multi-modal molecule structure–text model for text-based retrieval and editing. *Nature Machine Intelligence*, 5(12):1447–1457.

Zequn Liu, Wei Zhang, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Ming Zhang, and Tie-Yan Liu. 2023b. MolXPT: Wrapping molecules with text for generative pre-training. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1606–1616.

I Loshchilov. 2017. Decoupled weight decay regularization. In *International Conference on Learning Representations*.

OpenAI. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2024. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

Qizhi Pei, Lijun Wu, Kaiyuan Gao, Xiaozhuan Liang, Yin Fang, Jinhua Zhu, Shufang Xie, Tao Qin, and Rui Yan. 2024. BioT5+: Towards generalized biological understanding with IUPAC integration and multi-task tuning. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1216–1240.

Qizhi Pei, Wei Zhang, Jinhua Zhu, Kehan Wu, Kaiyuan Gao, Lijun Wu, Yingce Xia, and Rui Yan. 2023. BioT5: Enriching cross-modal integration in biology with chemical knowledge and natural language associations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1102–1123.

Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.

Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. 2022. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. 2022. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3):279–287.

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):513–530.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

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):862.

Jinlu Zhang, Yin Fang, Xin Shao, Huajun Chen, Ningyu Zhang, and Xiaohui Fan. 2024. The future of molecular studies through the lens of large language models. *Journal of Chemical Information and Modeling*, 64(3):563–566.

Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.

Alex Zhavoronkov. 2018. Artificial intelligence for drug discovery, biomarker development, and generation of novel chemistry. *Molecular Pharmaceutics*, 15(10):4311–4313.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv preprint arXiv:2403.13372*.

A Template Filling

We show the templates in Figure 5 for synthesizing instruction data.

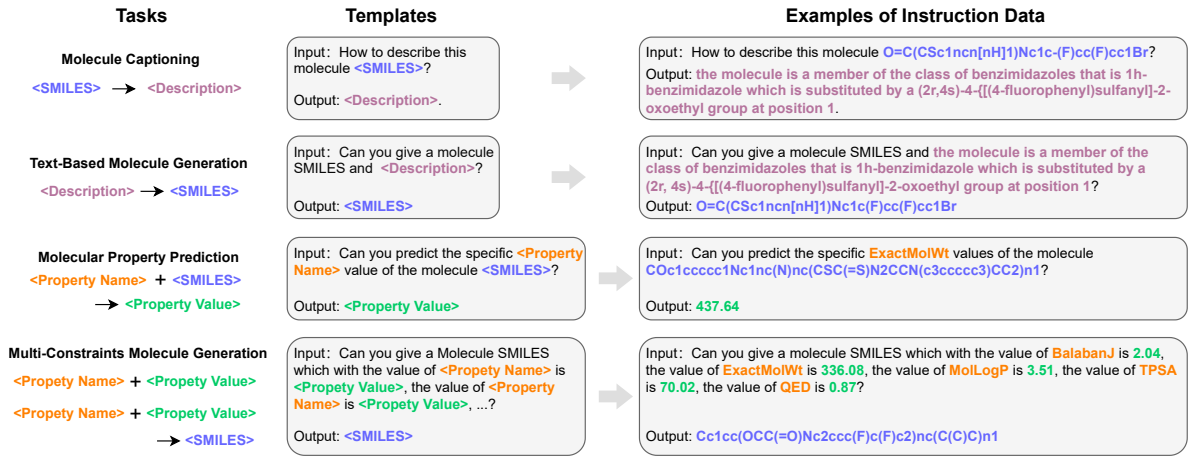

Figure 5: Examples of template filling with unstructured data according to four different downstream tasks for obtaining a variety of instruction data for supervised fine-tuning large language models.

B Details of Baselines

We describe the baseline models in our experiments as follows:

MolT5 (Edwards et al., 2022) is a framework for pre-training models on unlabeled text and molecular data. It introduces tasks like molecule captioning and generating molecules from text.

BioT5 (Pei et al., 2023) is a biology-focused pre-trained language model trained on diverse biological data, linking text with molecular and protein information.

BioT5+ (Pei et al., 2024) is a model optimized for biological research. It extends the BioT5 framework, enhancing the understanding and reasoning of biological texts and sequences, with notable success in molecule captioning and generation tasks.

MolXPT (Liu et al., 2023b) is a pre-trained language model for molecular science that enriches both text and molecular SMILES representations by replacing molecular names in the text with SMILES notation.

GIT-Mol (Liu et al., 2024) is a multi-modal LLM designed for molecular science, integrating graph, image, and text data. It performs well in tasks like molecule captioning, text-to-molecule generation, image recognition, and property prediction.

SPMM (Chang and Ye, 2024) is a multi-modal molecular pre-trained model that combines molecular structure information and biochemical properties by aligning two distinct features into a shared embedding space.

LLaMa3 (Touvron et al., 2023) is an open-source LLM, suitable for various NLP tasks such as summarization, question answering, and translation.

LLaMa3.1 (Dubey et al., 2024) is a series of updated open-source LLM based on LLaMa3, featuring a stronger parameter scale and higher performance.

Qwen2.5 (Yang et al., 2024) is an open-source large model that has been pre-trained on a dataset containing 18 trillion tokens. It has achieved significant improvements in overall capabilities and excels in a wide range of NLP tasks.

Mol-Instructions (Fang et al., 2023) is a natural language instruction dataset for biomolecules, designed to enhance the capabilities of large-scale pre-trained models in the biomolecular domain. This dataset combines biomolecules (such as proteins, DNA, RNA, etc.) with natural language instructions, supporting tasks such as molecule generation, molecule modification, and reaction prediction. We use the LLaMa3.1-8B model after SFT on this instruction dataset.

C Evaluation Metrics

We evaluated the quality of generated text using BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and ROUGE scores. These metrics evaluate the similarity between generated texts and reference descriptions, effectively quantifying the accuracy and diversity of the generated descriptions. For the text-based molecule generation task, we further use molecular fingerprints (FTS) (Cereto-Massagué et al., 2015) and validity measures to assess molecular similarity and validity, including Validity, Levenshtein (Levenshtein, 1966), MACCS FTS, Morgan FTS, and RDKit FTS (Landrum et al., 2013). For the task of molecular property prediction, we chose to use the

commonly used RMSE to measure the difference between the predicted values and the molecular property values calculated by RDKit for comparison. In the case of multi-constraint molecule generation, in addition to RMSE, we also employed R^2 to assess the accuracy of the generated molecules concerning multiple constrained properties.

D More Analysis on SFT Steps

The impact of SFT steps for text-based molecule generation and single MolLogP property prediction is shown in Figure 6 and Figure 7, respectively. Similar to two other tasks, we receive the best results through approximately 5∼6 epochs of SFT.

E Case Study

As for qualitative analysis, we show some case study for different tasks, directly comparing the response by LLaMa3.1 and our PEIT-LLM.

Table 6 shows the examples for molecule captioning. From the responses, we can see that the captions generated by our model are closer to the ground truth captions, whereas the original model often provides vague answers. This demonstrates that our model has a better understanding of molecules and can generate high-quality responses to complete molecular annotation tasks.

Table 7 shows the examples for text-based molecule generation. We can see that, compared with LLaMa3.1, the molecules generated by our model are more structurally reasonable and similar to the ground truth.

Table 8 shows the examples for molecular property prediction. LLaMa3.1 performs poorly in predicting specific molecular property and lacks the ability to accurately predict molecular property values. However, after instruction fine-tuning, our PEIT-LLM demonstrates excellent performance in predicting molecular property values, providing a strong foundation for its outstanding performance in multi-constraint molecule generation tasks.

Table 9 shows the examples for challenging multi-constraint molecule generation task. It can be observed that our model generated molecule with close properties according to multiple constraints. This indicates that the multi-task instruction data generated by our pre-trained model can enhance the capabilities of LLMs for such challenging task.

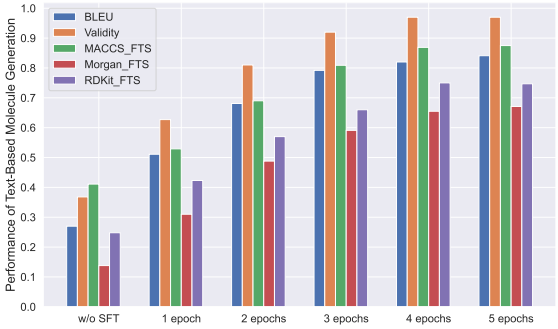

Figure 6: The impact of different amount of SFT steps for PEIT-LLM on text-based molecule generation task.

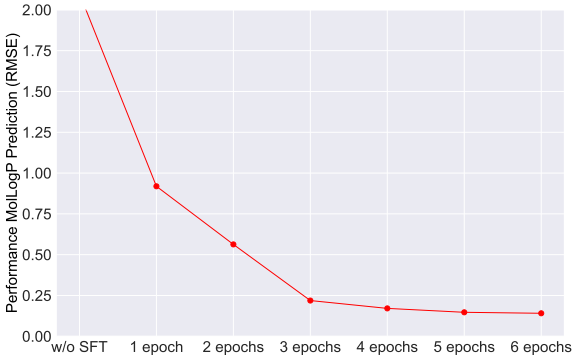

Figure 7: The impact of different amount of SFT steps for PEIT-LLM on MolLogP property prediction task.

Input for Molecule Captioning	LLaMa3.1 Output	Our PEIT-LLM Output	Ground Truth
1. How to describe this molecule CCCCCC1OC1C=CC(O)CC=CCC=CCCCC(=O)[O-]?	This is a complex organic molecule containing a cyclic ether, an alkene, an alcohol, a ketone, and an enolate anion.	The molecule is a prostaglandin carboxylic acid anion that is the conjugate base of prostaglandin F2alpha, obtained by deprotonation of the carboxy group; major species at pH 7.3. It is a conjugate base of a prostaglandin F2alpha.	The molecule is an epoxy(hydroxy)icosatrienoate that is the conjugate base of 11 hydroxy-(14R,15S)-epoxy-(5Z,8Z,12E)-icosatrienoic acid, obtained by deprotonation of the carboxy group; major species at pH 7.3. It is a conjugate base of an 11 hydroxy-(14R,15S)-epoxy-(5Z,8Z,12E)-icosatrienoic acid.
2. How to describe this molecule [Te]?	The molecule [Te] represents a tellurium atom in its elemental form.	The molecule is the stable isotope of tellurium with relative atomic mass 127.904, 7.00 atom percent natural abundance and nuclear spin 1/2.	The molecule is the stable isotope of tellurium with relative atomic mass 124.904425, 71.4 atom percent natural abundance and nuclear spin 1/2.
3. How to describe this molecule CN(N=O)C(N)=O?	This molecule is a nitrosoguanidine, characterized by a central carbon atom double-bonded to nitrogen and oxygen, with an additional nitrogen atom attached to the carbon.	The molecule is a nitrosamine that is urea in which one of the nitrogens is substituted by methyl and nitroso groups. It has a role as a carcinogenic agent, a mutagen, a teratogenic agent and an alkylating agent.	The molecule is a member of the class of N-nitrosoureas that is urea in which one of the nitrogens is substituted by methyl and nitroso groups. It has a role as a carcinogenic agent, a mutagen, a teratogenic agent and an alkylating agent.

Table 6: Examples of responses by LLaMa3.1 and our PEIT-LLM on molecule captioning.

Input for Text-based Molecule Generation	LLaMa3.1 Output	Our PEIT-LLM Output	Ground Truth
1. Can you give a molecule SMILES and the molecule is an amino dicarboxylic acid that is adipic acid in which one of the hydrogens at position 3 is replaced by an amino group. It is a beta-aminoacid an amino dicarboxylic acid and a gamma-amino acid. It derives from an adipic acid?	Cn1c(=O)c2c(nc(NCC3CCCO3)n2C C(=O)c2ccccc2)n(C)c1=O 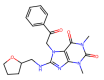	NC(CCCC(=O)O)C(=O)O 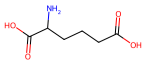	C(CC(=O)O)C(CC(=O)O)N 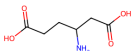
2. Can you give a molecule SMILES and the molecule is a methyl ketone that is decane in which the methylene hydrogens at position 2 are replaced by an oxo group. It has a role as a plant metabolite. It derives from a hydride of a decane?	CN(C)C1=CC=C2[NH1]C=3C(C([NH3 +1])=O)=CC(=O)C(=O)C=3OC2=C1 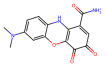	CCCCCCCC(=O)CC 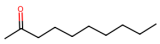	CCCCCCCCC(=O)C 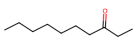
3. Can you give a molecule SMILES and the molecule is a member of the class of Nnitrosoureas that is urea in which one of the nitrogens is substituted by methyl and nitroso groups. It has a role as a carcinogenic agent, a mutagen, a teratogenic agent and an alkylating agent?	C1=CC(=O)OC2=CC(=C(C=C21)O [C@H]3[C@@H]([C@H]([C@@H] ([C@H](O3)CO)O)O)O)O 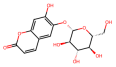	CCOP(=S)(OCC)OC1=NN(C=N1) C2=CC=CC=C2 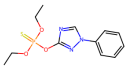	CCOP(=S)(OCC)OCC1=CC(=O) N(C2=CC=CC=C2)N1 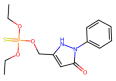

Table 7: Examples of responses by LLaMa3.1 and our PEIT-LLM for text-based molecule generation.

Input	LLaMa3.1 Output	Our PEIT-LLM Output	Ground Truth
1. Can you predict the specific MolLogP values of the molecule? CC1CC(C)CN(S(=O)(=O)c2ccc(C(=O)Nc3nnc(C4CC4)o3)cc2)C1	23.42	**2.65**	2.66
2. Can you predict the specific ExactMolWt values of the molecule? COc1ccc(NC(=O)Cn2ccn(-c3ccc(F)c(F)c3)c(=O)c2=O)cc1Cl	13.77	**404.94**	404.25
3. Can you predict the specific MolLogP values of the molecule? Cc1cc2c(cc1)C(=O)NC(C)C2	18.98	**2.53**	2.52

Table 8: Examples of responses by LLaMa3.1 and our PEIT-LLM for molecular property prediction.

Input for Multi-Constraint Molecule Generation	**LLaMa3.1 Output** ***(Verified Values by RDKit)***	**Our PEIT-LLM Output** ***(Verified Values by RDKit)***
1. Can you give a molecule SMILES which with the value of BalabanJ is 1.68, the value of ExactMolWt is 492.20, the value of MolLogP is 3.14, the value of TPSA is 138.76, the value of QED is 0.33?	NN=c1sc2ccccc2n1-c1ccccc1 *(BalabanJ is 2.42, ExactMolWt is 241.06, MolLogP is 2.46, TPSA is 43.30, QED is 0.51)*	CCOc1cc(C=NNC(=O)c2ccncc2)ccc1OS(=O)(=O)c1ccc(NC(C)=O)cc1 *(BalabanJ is 1.74, ExactMolWt is 482.12, MolLogP is 2.97, TPSA is 136.04, QED is 0.27)*
2. Can you give a molecule SMILES which with the value of BalabanJ is 1.33, the value of ExactMolWt is 467.20, the value of MolLogP is 3.54, the value of TPSA is 92.42, the value of QED is 0.52?	O=C(CN1CCN(c2ccc(Cl)cc2)CC1)Nc1ccc(F)cc1F *(BalabanJ is 1.49, ExactMolWt is 365.11, MolLogP is 3.37, TPSA is 35.58, QED is 0.90)*	O=C(COC(=O)c1ccc(S(=O)(=O)N2CCCc3ccccc32)cc1)Nc1ccc(F)cc1 *(BalabanJ is 1.39, ExactMolWt is 468.11, MolLogP is 3.76, TPSA is 92.78, QED is 0.55)*
3. Can you give a molecule SMILES which with the value of BalabanJ is 1.98, the value of ExactMolWt is 303.13, the value of MolLogP is 3.06, the value of TPSA is 64.11, the value of QED is 0.92?	O=C(NCc1cccc(F)c1)Nc1nnc(C2CC(O)C(CO)O2)s1 *(BalabanJ is 1.51, ExactMolWt is 368.09, MolLogP is 1.18, TPSA is 116.60, QED is 0.62)*	Cc1cccc(NS(=O)(=O)c2ccc3oc(C)c(C)c3c2)n1 *(BalabanJ is 2.11, ExactMolWt is 306.10, MolLogP is 2.88, TPSA is 67.43, QED is 0.90)*

Table 9: Examples of responses by LLaMa3.1 and our PEIT-LLM for multi-constraint molecule generation, and the verified property values of the output molecule are shown in the brackets.