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

Anonymous ACL submission

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 (**P**roperty **E**nhanced **I**nstruction **T**uning) 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, BioT5, MolCA and Text+Chem-T5 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, demonstrating the effectiveness of the PEIT framework for various molecular tasks.

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, 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., 2024a).

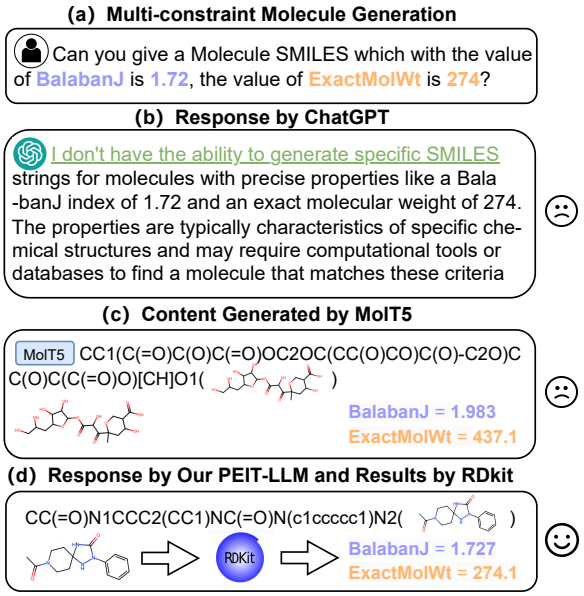

Figure 1: (a) An example of our proposed multi-constraint molecule generation task. (b) The response by ChatGPT. (c) The result generated by MolT5. (d) The response generated by the LLaMA3.1 model after applying our proposed property-enhanced instruction tuning, with the results validated by RDKit.

Despite their success, as illustrated in Figure 1 (a) and (b), LLMs still struggle with generating molecules under strict property constraints. Even specialized molecular translation models, such as MolT5 (Edwards et al., 2022), fail in these tasks (see Figure 1 (c)). While these models effectively capture relationships between molecular text and structure, they lack sufficient understanding of molecular properties, limiting their ability to incorporate property constraints in prompts. This shortcoming restricts their practical utility in appli-

cations like drug discovery (Zhavoronkov, 2018; Elton et al., 2019).

The challenges in addressing such tasks mainly lie in three aspects: (1) Existing studies have revealed the limitations of LLMs in understanding molecular representations (Grisoni, 2023), making it difficult to handle tasks requiring precise molecular property comprehension; (2) While there are some known SMILES-property pairing data, it often remains limited to predicting a single property and lacks datasets that cover 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 trimodal data pairs (Krenn et al., 2020); (3) To our knowledge, there are no suitable datasets or evaluation methods exist for multi-constraint molecule generation using LLMs, which challenges the standardization and assessment of such tasks (Jin et al., 2018; Elton et al., 2019).

To address these challenges, we propose a framework called PEIT (**P**roperty **E**nhanced **I**nstruction **T**uning) 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 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 2. Specifically, it consists of two components: (1) We pretrain 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 finetune 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.

Experimental results demonstrate that our pretrained 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), MolCA (Liu et al., 2023c), and Text+Chem-T5 (Christofidellis et al., 2023). Additionally, PEIT-LLM based on LLaMa3.1-8B (Dubey et al., 2024) exhibits superior performance compared to specialized models Mol-Instructions (Fang et al., 2023) and general-purpose LLMs including LLaMa3 (Dubey et al., 2024) and Qwen2.5 (Yang et al., 2024) in molecular property prediction and our newly proposed multi-constraint molecule generation tasks.

2 Related Work

Molecule Generation. Molecule generation tasks mainly fall into two categories: (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

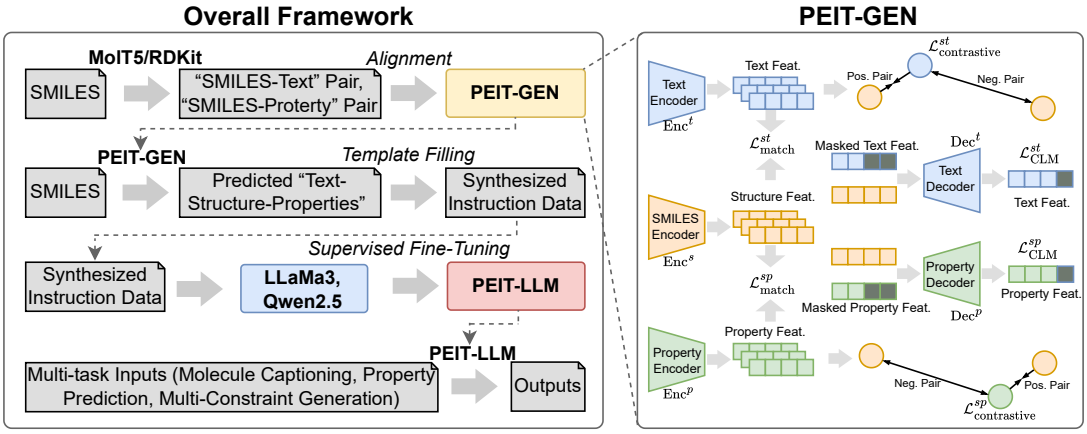

Figure 2: Left: Overall PEIT framework. We first pre-train the PEIT-GEN and construct instruction data via template filling. Then we fine-tune the open-source LLMs through instruction tuning, the resulting PEIT-LLM is used for multi-task molecule generation. Right: The process of PEIT-GEN pre-training, see details in Section 3.2.

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. Constructing specialized instruction datasets is an effective way to enable LLMs to better perform molecular-related tasks. For instance, Mol-Instructions (Fang et al., 2023) provides a large-scale biomolecular dataset tailored for LLMs, covering a variety of instructions involving small molecules, proteins, and biomolecular texts. ChemDFM (Zhao et al., 2025) advances this paradigm by creating a broader dataset spanning molecular structures, reactions, and properties. Its two-stage training—domain pretraining followed by instruction tuning—enhances the model's chemical understanding and reasoning capabilities. More recently, GeLLM3O (Dey et al., 2025) introduced MuMOInstruct, a high-quality dataset focused on multi-property molecule optimization, demonstrating strong generalization across diverse tasks. Despite these advances, generating reliable and scalable molecular instruction data remains a key challenge, particularly for open-source models.

3 Method

3.1 Overview of PEIT Framework

The overview of PEIT framework is shown in Figure 2 (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 2. 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. Following SPMM (Chang and Ye, 2024), 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:

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

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

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:

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

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

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:

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

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

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.

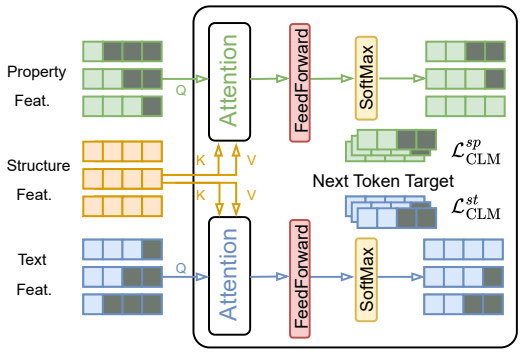

Figure 3: The cross-modal causal language modeling.

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:

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

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

Cross-modal Causal Language Modeling. To further strengthen the model's capability in molecule captioning, we employ the causal language modeling (CLM) to enhance the model performance on text generation. Specifically, we design decoders to generate subsequent property and textual description sequences, under the guidance of SMILES features through cross-attention as show in Figure 3.

By introducing the SMILES features in attention layers for CLM training, the cross-modal CLM loss $\mathcal{L}^{st}_{\text{CLM}}$ and $\mathcal{L}^{sp}_{\text{CLM}}$ are computed as follows:

$$\mathcal{L}^{st}_{\text{CLM}} = -\sum_{i=1}^{N}\sum_{j=1}^{n}\log \text{Prob}\left(w^{(i)}_j \mid \text{Dec}^t(\tilde{\mathbf{w}}^{(i)}_{:j}); \theta_t\right), \quad (9)$$

$$\mathcal{L}^{sp}_{\text{CLM}} = -\sum_{i=1}^{N}\sum_{j=1}^{n}\log \text{Prob}\left(w^{(i)}_j \mid \text{Dec}^p(\tilde{\mathbf{w}}^{(i)}_{:j}); \theta_p\right), \quad (10)$$

where Prob is the conditional probability to predict the word $w^{(i)}_j$ in the vocabulary, N is the total number of samples, n is the index of current words in each sample, $\tilde{\mathbf{w}}^{(i)}_{:j}$ is the sequence from begin to the j-th word in the i-th sample, θ_t and θ_p are the trainable parameters in two decoders.

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

$$\mathcal{L} = \mathcal{L}^{st}_{\text{match}} + \mathcal{L}^{sp}_{\text{match}} + \alpha\mathcal{L}^{st}_{\text{contrastive}} + \alpha\mathcal{L}^{sp}_{\text{contrastive}} + \beta\mathcal{L}^{st}_{\text{CLM}} + \beta\mathcal{L}^{sp}_{\text{CLM}}, \quad (11)$$

where we follow SPMM (Chang and Ye, 2024) to use parameters α and β (α:β=1:5) for balancing 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, which are stored in CSV files containing text, molecular structures, and information on 53 molecular biochemical properties. 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 7 in Appendix A. For text-based molecule generation as example, we fix a general question format and then extract molecular descriptions from unstructured data to fill the 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.

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) iven the vast number of molecular attributes and the complex combinations thereof, analyzing their impact on results across different molecular counts poses a significant computational challenge due to resource limitations. To address this, we selected representative ADMET (Fu et al., 2024) properties—BalabanJ, MolLogP, ExactMolWt, QED, and TPSA—that capture molecular topology, electronic characteristics, and steric effects, while exhibiting low mutual correlation. Based on template filling, the predicted multi-property values are used to construct data for multi-constraint molecule generation. We employ instruction tuning to guide the LLM in generating molecules, and use RDKit (Landrum et al., 2013) to calculate the actual property values. RMSE and R^2 are then used to compare these values against the constraints, enabling a systematic evaluation of the LLM's performance in multi-constraint molecule generation tasks.

Model	MC	TBMG	MPP	MCMG
MolT5	✓	✓	✗	✗
BioT5	✓	✓	✗	✗
MolXPT	✓	✓	✗	✗
Git-Mol	✓	✓	✗	✗
SPMM	✗	✗	✓	✗
MolCA	✓	✓	✗	✗
Text+Chem-T5	✓	✓	✗	✗
BioMedGPT	✓	✗	✗	✗
InstructMol-GS	✓	✗	✗	✗
MolReGPT	✓	✓	✗	✗
Mol-Instructions	✓	✓	✓ (poor)	✓ (poor)
LLaMa, Qwen	✓ (limited)	✓ (poor)	✓ (poor)	✓ (poor)
PEIT-LLM (Ours)	✓	✓	✓	✓

Table 1: Comparing PEIT-LLM with biomolecular models and LLMs on molecular-related tasks. MC: Molecule Captioning. TBMG: Text-Based Molecule Generation. MPP: Molecular Property Prediction. MCMG: 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 Large Language Models

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

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
MolXPT (Liu et al., 2023b)†	30M	0.594	0.505	0.626	0.660	0.511	0.597
$\text{MolCA}_{\text{w/ Galac}}$ (Liu et al., 2023c)	2.3M	0.616	0.524	0.639	0.674	0.533	0.615
$\text{Text+Chem-T5}_{\text{augm}}$ (Christofidellis et al., 2023)	11.5M	0.625	0.529	0.648	0.682	0.543	0.622
GIT-Mol (Liu et al., 2024)†	4.8M	0.352	0.263	0.533	0.575	0.485	0.560
PEIT-GEN (Ours)	**0.48M**	0.598	0.534	**0.676**	**0.700**	**0.582**	**0.653**

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

4 Experiments

4.1 Experimental Setup

Dataset. For pre-training PEIT-GEN, we extract 480k molecular SMILES from the ZINC dataset (Irwin et al., 2012) and generate SMILES-text pairs using MolT5 (Edwards et al., 2022). We also compute 53 biochemical properties per molecule using RDKit, forming 480k "text-SMILES-properties" triplets for training.

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.

To evaluate PEIT-GEN and PEIT-LLM, we follow MolT5 by using CHEBI-20 (Edwards et al., 2021) and MoleculeNet dataset (Wu et al., 2018), with the standard splition into training, validation, and test sets with an 8:1:1 ratio. All property values are verified via RDKit. See details in Appendix B.

Baseline Models. We compare our model, PEIT-GEN and PEIT-LLM, against three types of baselines as follows: *Baselines on molecule caption* such as MolT5 (Edwards et al., 2022), BioT5 (Pei et al., 2023), MolCA (Liu et al., 2023c), Text+Chem-T5 (Christofidellis et al., 2023), GIT-Mol (Liu et al., 2024). *Baselines on moleucular property prediction* such as SPMM (Chang and Ye, 2024), D-MPNN (Yang et al., 2019), PretrainGNN (Hu et al., 2019), GROVER (Rong et al., 2020), ChemRL-GEM (Fang et al., 2022). *Baselines of LLMs* such as LLaMa3 (Touvron et al., 2023), Qwen2.5 (Yang et al., 2024), GPT3.5-turbo (OpenAI, 2023) , Mol-Instructions (Fang et al., 2023), InstructMol-GS (Cao et al., 2023), BioMedGPT (Zhang et al., 2024b), ChemDFM (Zhao et al., 2025), ChatGLM (GLM et al., 2024), Galactica (Taylor et al., 2022), Vicuna (Chiang et al., 2023). Details of these baselines and evaluation metric are in Appendix C and D, respectively.

Model	BBBP	BACE	Clintox	SIDER
D-MPNN (Yang et al., 2019)	71.0±0.3	80.9±0.6	90.6±0.6	57.0±0.7
N-GramRF (Liu et al., 2019)	69.7±0.6	77.9±1.5	77.5±4.0	66.8±0.7
N-GramXGB (Liu et al., 2019)	69.1±0.8	79.1±1.3	87.5±2.7	65.5±0.7
PretrainGNN (Hu et al., 2019)	68.7±1.3	84.5±0.7	72.6±1.5	62.7±0.8
$\text{GROVER}_{\text{large}}$ (Rong et al., 2020)	69.5±0.1	81.0±1.4	76.2±3.7	65.4±0.1
ChemRL-GEM (Fang et al., 2022)	72.4±0.4	85.6±1.1	90.1±1.3	**67.2±0.4**
ChemBERTa (Ahmad et al., 2022)†	72.8	79.9	56.3	-
MolFormer (Ross et al., 2022)	73.6±0.8	**86.3±0.6**	91.2±1.4	65.5±0.2
SPMM (Chang and Ye, 2024)	**74.1±0.6**	82.9±0.3	90.7±0.5	63.6±0.5
PEIT-GEN (Ours)	73.6±0.7	81.6±0.5	**91.2±0.7**	62.7±0.9

Table 3: Results on MoleculeNet dataset for 2-way property prediction. †: The standard deviation and the results on SIDER are not reported in literature.

Implementation Details. We pre-train PEIT-GEN for 20 epochs using a batch size of 16, temperature $\tau = 0.07$, and momentum 0.995 with the AdamW optimizer (Loshchilov, 2017). Fine-tuning is then performed on the CHEBI-20 training set for 200 epochs with a learning rate of 5e-4. For supervised fine-tuning of PEIT-LLM, we use the LLaMa-Factory (Zheng et al., 2024) framework with LoRA (Hu et al., 2022) for 6 epochs, a batch size of 3, and learning rate of 5e-5. The total parameter count of the three encoders and two decoders is 533.8M, with an additional 2.2M parameters for other components. Experiments are conducted on NVIDIA 4090 GPUs with 24GB memory.

4.2 Comparing PEIT-GEN with Pre-trained Biomolecular Models

Molecule Captioning. Results on molecule captioning using CHEBI-20 dataset 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-4. Compared to BioT5 which performs the best in BLEU, our approach

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
GPT-3.5-turbo (OpenAI, 2023)	N/A†	0.103	0.050	0.161	0.261	0.088	0.204
Mol-Instructions (Fang et al., 2023)	8B	0.217	0.143	0.254	0.337	0.196	0.291
BioMedGPT (Zhang et al., 2024b)	10B	0.234	0.141	0.308	0.386	0.206	0.332
InstructMol-GS (Cao et al., 2023)	7B	0.475	0.371	0.509	0.566	0.394	0.502
MolReGPT (Li et al., 2024)	N/A†	**0.565**	**0.482**	**0.585**	**0.623**	**0.450**	**0.543**
ChemDFM (Zhao et al., 2025)	13B	0.321	0.265	0.402	0.490	0.374	0.483
PEIT-LLM-Qwen2.5 (Ours)	7B	0.422	0.314	0.468	0.535	0.361	0.477
PEIT-LLM-LLaMa3.1 (Ours)	8B	0.461	0.356	0.502	0.569	0.396	0.505

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
GPT-3.5-turbo (OpenAI, 2023)	N/A†	0.489	0.802	52.130	0.705	0.367	0.462
Mol-Instructions (Fang et al., 2023)	8B	0.345	**1.000**	41.367	0.412	0.147	0.231
MolReGPT (Li et al., 2024)	N/A†	0.790	0.887	24.910	0.847	0.624	0.708
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 CHEBI-20 dataset for molecule captioning (top) and text-based molecule generation (bottom) tasks. †: MolReGPT is based on closed-source ChatGPT-3.5 and its parameter size remains unknown.

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.

Molecular Property Prediction. We evaluate the generalization ability of PEIT-GEN on the MoleculeNet benchmark (Wu et al., 2018) using four widely adopted tasks. As shown in Table 3, PEIT-GEN outperforms specialized models such as MolFormer (Ross et al., 2022) and ChemRL-GEM (Fang et al., 2022) on the Clintox dataset. Despite using less pre-training data, it remains competitive on other subsets. To further demonstrate its predictive strength across 53 molecular properties, we present a relative difference analysis in Appendix E, highlighting PEIT-GEN's strong generalization in property prediction.

4.3 Comparing PEIT-LLM with LLMs

Molecule Captioning. As shown in the top of Table 4, the comparison results show that our model outperforms general-purpose Qwen-2.5 and LLaMa3.1 as well as Mol-Instructions and BioMedGPT, which were trained using a biochemical information instruction dataset for SFT. PEIT-LLM achieved the second-best performance on the ROUGE metric and demonstrated competitive results compared to InstructMol-GS, which was trained solely on the CHEBI-20 dataset and has a similar parameter scale as our base model.

Text-based Molecule Generation. Results on the CHEBI-20 test set are presented at the bottom of Table 4. PEIT-LLM outperforms all baselines

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.

on numerical metrics, including BLEU, Levenshtein Distance, and fingerprint similarities based on MACCS, Morgan, and RDKit. Although Mol-Instructions achieves the highest Validity score, the results demonstrate that PEIT-LLM, after multi-task instruction fine-tuning, effectively captures key molecular structures and their corresponding textual representations. A case study in Table 9 of Appendix F further supports these findings and indirectly validates the quality of the instruction data generated by PEIT-GEN.

Molecular Property Prediction. For single-property prediction, due to the large number of available properties, we select two representative examples: ExactMolWt, which typically has large numerical values (100–1000), and MolLogP, with smaller values (–5 to 10), as shown in Table 5. The results show that PEIT-LLM consistently outperforms other LLMs in predicting these biochemical properties, demonstrating strong sensitivity and adaptability to molecular property scales. This highlights the effectiveness of multi-task SFT in enhancing LLMs' understanding of molecular char-

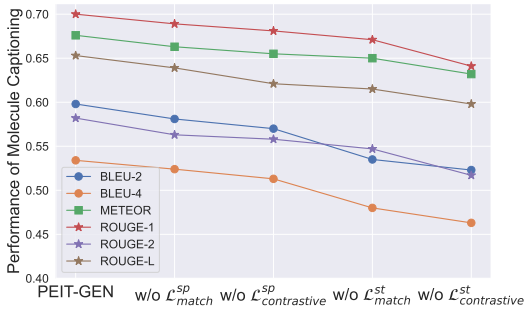

Figure 4: Ablation study on PEIT-GEN pre-training objectives $\mathcal{L}^{sp}_{\text{match}}$, $\mathcal{L}^{st}_{\text{match}}$, $\mathcal{L}^{sp}_{\text{contrastive}}$, and $\mathcal{L}^{st}_{\text{contrastive}}$.

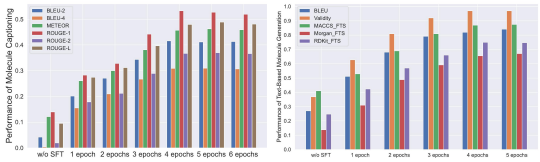

Figure 5: The impact of different amount of SFT steps on molecule captioning (left) and generation (right).

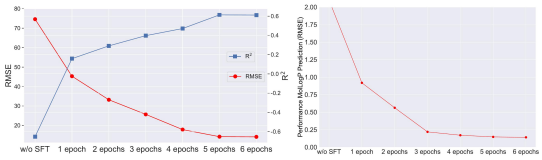

Figure 6: The impact of different amount of SFT steps on generation (left) and prediction (right) tasks.

acteristics and further validates the quality and reliability of our molecular property instruction dataset. A case study is provided in Table 10 of Appendix F for further illustration.

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 11 of Appendix F 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. The model gain strong molecular understanding capabilities through property enhanced instruction tuning.

4.4 Analyses

Ablation Study. Figure 4 presents an ablation study of the cross-modal matching loss $\mathcal{L}_{\text{match}}$ and cross-modal contrastive loss $\mathcal{L}_{\text{contrastive}}$ in the PEIT-GEN model for the molecule captioning task ($\mathcal{L}^{st}_{\text{CLM}}$ and $\mathcal{L}^{sp}_{\text{CLM}}$ are necessary for generation via decoders, thus we do not consider them in ablation study). By removing these training objectives, the performance degradation across all metrics. This demonstrates that both $\mathcal{L}_{\text{match}}$ and $\mathcal{L}_{\text{contrastive}}$ are helpful in cross-modal feature alignment, thereby enhancing the performance of molecule captioning.

Model	BLEU↑	METEOR↑	ROUGE-2↑	ROUGE-L↑
Galactica-6.7B (Taylor et al., 2022)	0.008	0.065	0.015	0.063
MolT5-248M (Edwards et al., 2022)	0.001	0.033	0.001	0.034
Vicuna-7B (Chiang et al., 2023)	0.011	0.168	0.055	0.130
Text+Chem T5-223M (Christofidellis et al., 2023)	0.036	0.139	0.075	0.119
ChatGLM-6B (GLM et al., 2024)	0.011	0.105	0.066	0.148
LLaMa3.1-8B (Dubey et al., 2024)	0.014	0.184	0.066	0.148
Qwen2.5-7B (Yang et al., 2024)	0.009	0.169	0.047	0.119
PEIT-LLM-Qwen2.5-7B (Ours)	0.051	0.208	0.121	0.178
PEIT-LLM-LLaMa3.1-8B (Ours)	0.053	0.215	0.125	0.184
Mol-Instructions-8B (Fang et al., 2023)	0.143	0.254	0.196	0.291

Table 6: Out-of-distribution results on the molecule captioning task using Mol-Instructions (Fang et al., 2023) evaluation set. Mol-Instructions denotes a fully baseline trained with LLaMA3.1-8B using the entire training instructions, serving as an upper bound for SFT models.

Impact of SFT steps. Figure 5 and Figure 6 illustrate the outcomes of PEIT-LLM across various tasks with different SFT steps. We observe that the performance consistently improves during the initial epochs for all tasks, indicating that the instructional data is beneficial for each, where the performance tends to plateau around epochs 5-6.

Out-of-Distribution Evaluation. To further evaluate the molecular understanding of PEIT-LLM on unseen data, we tested it on the Mol-Instructions (Fang et al., 2023) test set, without using the full instructions for pre-training and fine-tuning PEIT-LLMs. As shown in Table 6, PEIT-LLM outperforms all general-purpose LLMs as well as smaller domain-specific models highlighting the strong generalization ability of PEIT-LLM across more diverse molecular instruction tasks.

5 Conclusion

We propose PEIT, a framework that enables LLMs to perceive multi-modal features for multi-task molecule generation. PEIT aligns molecular structures, textual descriptions, and biochemical properties through multi-modal representation learning. It leverages templates to synthesize diverse, task-specific instruction data for LLMs. We also introduce a challenging multi-constraint molecule generation task, which requires generating novel molecules that satisfy multiple property constraints. Results show that PEIT outperforms various biomolecular models and LLMs on captioning, generation, and property prediction tasks.

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

Walid Ahmad, Elana Simon, Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. 2022. Chemberta-2: Towards chemical foundation models. *arXiv preprint arXiv:2209.01712*.

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.

Leo Breiman. 2001. Random forests. *Machine learning*, 45:5–32.

He Cao, Zijing Liu, Xingyu Lu, Yuan Yao, and Yu Li. 2023. Instructmol: Multi-modal integration for building a versatile and reliable molecular assistant in drug discovery. *Preprint*, arXiv:2311.16208.

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.

Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.

Zhihao Chiang, Zhuohan Zhu, Xinyang Zeng, Yian Xu, Zi Lin, Hao Li, Lianmin Zhang, Shreyas Jain, Rohan Taori, Ishaan Dave, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. *Preprint*, arXiv:2304.05334.

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.

Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo Manica. 2023. Unifying molecular and textual representations via multi-task language modelling. In *Proceedings of the 40th International Conference on Machine Learning*, pages 6140–6157.

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.

Vishal Dey, Xiao Hu, and Xia Ning. 2025. GeLLM^{3O}: Generalizing large language models for multi-property molecule optimization. *arXiv preprint arXiv:2502.13398*.

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.

Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang, Hua Wu, and Haifeng Wang. 2022. Geometry-enhanced

molecular representation learning for property prediction. *Nature Machine Intelligence*, 4(2):127–134.

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*.

Li Fu, Shaohua Shi, Jiacai Yi, Ningning Wang, Yuanhang He, Zhenxing Wu, Jinfu Peng, Youchao Deng, Wenxuan Wang, Chengkun Wu, et al. 2024. Admetlab 3.0: an updated comprehensive online admet prediction platform enhanced with broader coverage, improved performance, api functionality and decision support. *Nucleic acids research*, 52(W1):W422–W431.

Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.

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*.

Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2019. Strategies for pre-training graph neural networks. *arXiv preprint arXiv:1905.12265*.

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*.

Jiatong Li, Yunqing Liu, Wenqi Fan, Xiao-Yong Wei, Hui Liu, Jiliang Tang, and Qing Li. 2024. Empowering molecule discovery for molecule-caption translation with large language models: A chatgpt perspective. *IEEE transactions on knowledge and data engineering*.

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, Mehmet F Demirel, and Yingyu Liang. 2019. N-gram graph: Simple unsupervised representation for graphs, with applications to molecules. *Advances in neural information processing systems*, 32.

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.

Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. 2023c. Molca: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In *EMNLP*.

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, 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*.

Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. 2020. Self-supervised graph transformer on large-scale molecular data. *Advances in neural information processing systems*, 33:12559–12571.

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.

Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*.

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*.

K. Yang et al. 2019. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388.

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. 2024a. The future of molecular studies through the lens of large language models. *Journal of Chemical Information and Modeling*, 64(3):563–566.

Kai Zhang, Rong Zhou, Eashan Adhikarla, Zhiling Yan, Yixin Liu, Jun Yu, Zhengliang Liu, Xun Chen, Brian D Davison, Hui Ren, et al. 2024b. A generalist vision–language foundation model for diverse biomedical tasks. *Nature Medicine*, pages 1–13.

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*.

Zihan Zhao, Da Ma, Lu Chen, Liangtai Sun, Zihao Li, Yi Xia, Bo Chen, Hongshen Xu, Zichen Zhu, Su Zhu, et al. 2025. Developing chemdfm as a large language foundation model for chemistry. *Cell Reports Physical Science*, 6(4).

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 7 for synthesizing instruction data.

B Details of Property Prediction Tasks

Following SPMM (Chang and Ye, 2024) , we adopt four commonly-used binary property prediction tasks to evaluate the performance of PEIT-GEN, including BBBP, BACE, Clintox, and SIDER dataset. The BBBP dataset contains 2,050 molecular samples and aims to predict whether these molecules can cross the blood-brain barrier. The BACE dataset includes 1,513 molecular samples and is used to predict whether a molecule can inhibit the activity of the BACE1 enzyme. The Clintox dataset contains 1,478 molecular samples and is primarily used to predict the toxicity of compounds. The SIDER dataset consists of 1,427 drug samples and is used to predict whether a drug will cause specific side effects. Specifically, we use scaffold splitting and each dataset is divided into a training set, validation set, and test set in a ratio of 8:1:1, respectively.

C Details of Baselines

We compare our model against a variety of baselines which can be categorized as follows:

Baselines on molecule captioning task:

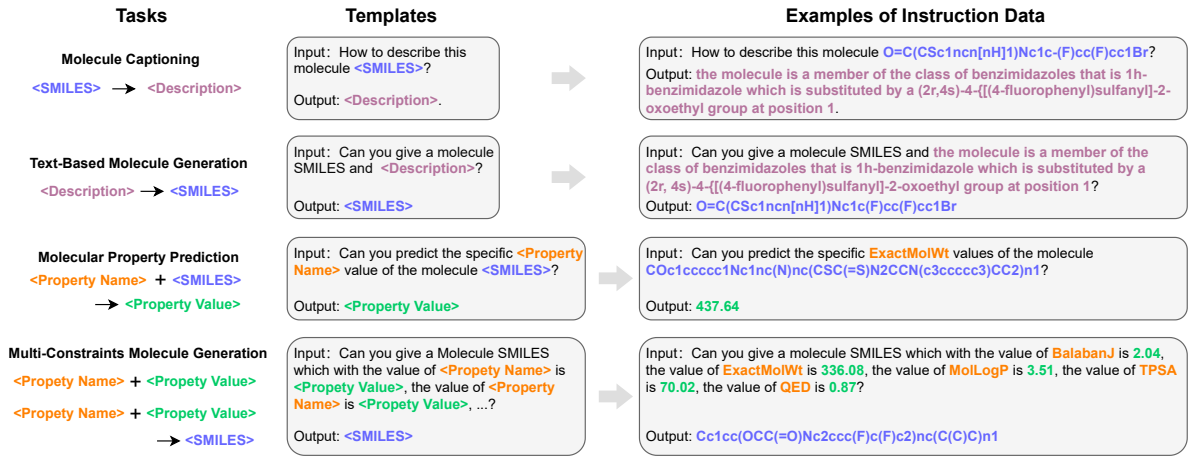

Figure 7: Examples of template filling with unstructured data according to four different downstream tasks for obtaining a variety of instruction data for supervised fine-tuning large language models.

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.

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.

MolCA (Liu et al., 2023c) is a model that combines molecular graphs with textual descriptions, excelling in molecular representation learning, cross-modal reasoning, and tasks such as property prediction, generation, and interaction.

Text+Chem-T5 (Christofidellis et al., 2023) is a multimodal model based on the T5 architecture, specifically designed for joint chemistry-text tasks. By integrating chemical data with natural language text, it enhances performance in chemical text understanding, molecular property prediction, and reaction generation tasks.

Baselines on molecular 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.

D-MPNN (Yang et al., 2019) D-MPNN is specifically designed for processing molecular graph data. It efficiently captures atomic interactions and chemical bond information through a directed message-passing mechanism, providing strong support for molecular property prediction.

N-GramRF (Liu et al., 2019) extracts N-Gram features from molecular sequences and integrates them with a Random Forest (Breiman, 2001) model to capture local structural information of molecules. It is suitable for molecular property prediction tasks, offering strong robustness and easy implementation.

N-GramXGB (Liu et al., 2019) also utilizes N-Gram features but employs the XGBoost (Chen and Guestrin, 2016) model for prediction. It efficiently handles high-dimensional data and captures nonlinear relationships, often outperforming Random Forest in predictive performance.

PretrainGNN (Hu et al., 2019) performs pre-training on molecular graph-structured data through self-supervised learning tasks, thereby learning universal representations of nodes and edges within the graph. This significantly enhances the model's performance in molecular property prediction tasks.

GROVER (Rong et al., 2020) leverages multiple self-supervised learning tasks to learn universal representations of atoms and bonds in molecular structures, significantly enhancing performance in downstream tasks such as molecular property prediction and drug discovery.

ChemRL-GEM (Fang et al., 2022) employs Graph Neural Networks (GNNs) to learn the embedding

representations of molecular graphs and utilizes reinforcement learning to optimize these representations, thereby better accomplishing tasks such as molecular property prediction and molecular generation.

ChemBERTa (Ahmad et al., 2022) is pre-trained on a large-scale chemical literature and biomedical corpora, learning linguistic features specific to the chemistry and biomedical domains. This enables it to excel in tasks such as molecular property prediction, drug discovery, and biomedical text mining.

MolFormer (Ross et al., 2022) captures global atomic interactions within molecules using self-attention and learns universal molecular representations through pretraining on large-scale datasets, demonstrating strong performance in property prediction and molecular generation tasks.

Baselines of LLMs:

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.

GPT-3.5 Turbo (OpenAI, 2023) is an advanced large language model developed by OpenAI, optimized for efficient inference and versatile natural language understanding and generation tasks. Built upon the transformer architecture, GPT-3.5 Turbo demonstrates strong performance across a wide range of NLP benchmarks, including text completion, summarization, translation, and dialogue systems. Its design balances high accuracy with reduced computational cost, making it suitable for scalable real-world applications.

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.

BioMedGPT (Zhang et al., 2024b) is a multimodal pre-trained model for the biomedical field, leveraging self-supervised learning and cross-modal alignment to learn universal representations from large-scale data, excelling in text understanding, medical image analysis, and molecular property prediction.

InstructMol-GS (Cao et al., 2023) is an instruction-tuned molecular generation model that maps natural language to molecular structures, enabling targeted molecule design and demonstrating strong generative capabilities in drug discovery and materials science.

MolReGPT (Li et al., 2024) is a molecule-text translation framework based on LLMs. It utilizes a molecular similarity retrieval mechanism to select examples, enabling efficient molecule generation and understanding without fine-tuning.

ChemDFM (Zhao et al., 2025) is a large language foundation model for the field of chemistry. Trained on 34 billion tokens from chemical literature and 2.7 million instructions, it demonstrates strong capabilities in understanding and reasoning about chemical knowledge. It supports tasks such as molecule recognition, design, property prediction, and reaction analysis, outperforming many open-source large language models.

ChatGLM (GLM et al., 2024) is an open-source bilingual large language model optimized for Chinese and English. It supports instruction tuning and multi-round dialogue, making it adaptable to domain-specific tasks such as molecular captioning and property prediction with appropriate prompting.

Galactica (Taylor et al., 2022) is a large language model pretrained on scientific texts, including papers, molecules, and protein sequences. Designed to assist scientific reasoning and knowledge retrieval, it supports molecule-related tasks to a limited extent through its exposure to structured scientific data during pretraining.

Vicuna (Chiang et al., 2023) is an open-source language model fine-tuned from LLaMA using user-shared conversations. It focuses on improving instruction-following and dialogue capabilities, and can be adapted to domain-specific tasks through fine-tuning, despite lacking scientific domain pretraining.

D Evaluation Metrics

We evaluated the quality of generated text using BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and ROUGE scores. These

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

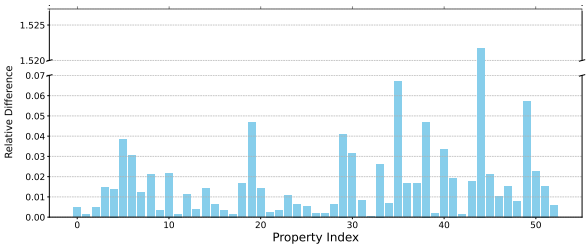

Figure 8: The relative difference represent the variation between the PEIT-GEN predicted values and the actual values for 53 distinct molecular properties.

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, for the experiments on MoleculeNet, we use AUC-ROC to evaluate the accuracy for property prediction tasks. 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.

E Molecular Property Prediction

Following SPMM (Chang and Ye, 2024), we further compare PEIT-GEN with SPMM on external dataset. The comparison result on molecular property prediction is shown in Table 7. Specifically, we randomly sample 1,000 molecules from the ZINC dataset which are not included in the training set. Compared to SPMM that is specifically designed for property prediction, PEIT-GEN achieves comparable performance while using only one-third of the data size across three modalities. We found that PEIT-GEN outperformed SPMM in terms of RMSE, while SPMM was slightly ahead by 0.11% on R^2 metric. These results demonstrate that PEIT-GEN can generate high-quality biochemical properties of molecules, highlighting the critical role of high-quality multi-modal data in advancing molecular understanding tasks. To further illustrate the predictive performance of PEIT-GEN on the 53 molecular properties, we calculated the relative difference of these properties, as shown in Figure 8.

F 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 8 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 9 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 10 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 11 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.

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 8: 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 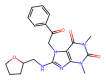	NC(CCCC(=O)O)C(=O)O 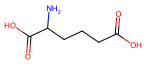	C(CC(=O)O)C(CC(=O)O)N 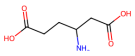
2. Can you give a molecule SMILES and the molecule is a methyl ketone that is decane in which the methylene hydrogens at position 2 are replaced by an oxo group. It has a role as a plant metabolite. It derives from a hydride of a decane?	CN(C)C1=CC=C2[NH1]C=3C(C([NH3 +1])=O)=CC(=O)C(=O)C=3OC2=C1 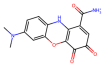	CCCCCCCC(=O)CC 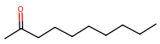	CCCCCCCCC(=O)C 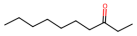
3. Can you give a molecule SMILES and the molecule is a member of the class of Nnitrosoureas that is urea in which one of the nitrogens is substituted by methyl and nitroso groups. It has a role as a carcinogenic agent, a mutagen, a teratogenic agent and an alkylating agent?	C1=CC(=O)OC2=CC(=C(C=C21)O [C@H]3[C@@H]([C@H]([C@@H] ([C@H](O3)CO)O)O)O)O 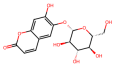	CCOP(=S)(OCC)OC1=NN(C=N1) C2=CC=CC=C2 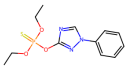	CCOP(=S)(OCC)OCC1=CC(=O) N(C2=CC=CC=C2)N1 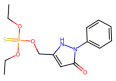

Table 9: Examples of responses by LLaMa3.1 and our PEIT-LLM for text-based molecule generation.

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 10: 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 11: 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.