

Improving Large Molecular Language Model via Relation-aware Multimodal Collaboration

Jinyoung Park*, Minseong Bae*, Jeehye Na*, Hyunwoo J. Kim[†]

Korea Advanced Institute of Science and Technology
{jinyoung.park, bms2002, jeehyena, hyunwoojkim}@kaist.ac.kr

Abstract

Large language models (LLMs) have demonstrated their instruction-following capabilities and achieved powerful performance on various tasks. Inspired by their success, recent works in the molecular domain have led to the development of large molecular language models (LMLMs) that integrate 1D molecular strings or 2D molecular graphs into the language models. However, existing LMLMs often suffer from hallucination and limited robustness, largely due to inadequate integration of diverse molecular modalities such as 1D sequences, 2D molecular graphs, and 3D conformations. To address these limitations, we propose CoLLaMo, a large language model-based molecular assistant equipped with a multi-level molecular modality-collaborative projector. The relation-aware modality-collaborative attention mechanism in the projector facilitates fine-grained and relation-guided information exchange between atoms by incorporating 2D structural and 3D spatial relations. Furthermore, we present a molecule-centric new automatic measurement, including a hallucination assessment metric and GPT-based caption quality evaluation to address the limitations of token-based generic evaluation metrics (*i.e.*, BLEU) widely used in assessing molecular comprehension of LMLMs. Our extensive experiments demonstrate that our CoLLaMo enhances the molecular modality generalization capabilities of LMLMs, achieving the best performance on multiple tasks, including molecule captioning, computed property QA, descriptive property QA, motif counting, and IUPAC name prediction.

Code — <https://github.com/mlvlab/CoLLaMo>

1 Introduction

Large language models (LLMs) (Touvron et al. 2023; Achiam et al. 2023; Team et al. 2023; Liu et al. 2023a; Dai et al. 2023b) have shown their significant advancements in complex reasoning tasks and visual tasks. Inspired by the progress of the LLMs, growing efforts have been made to adapt language modeling to diverse domains such as table understanding (Zhang et al. 2024; Chen et al. 2024; Kim et al. 2026), arithmetic reasoning (Shao et al. 2024;

Yang et al. 2024; Chu, Lee, and Kim 2025), and scientific domains, particularly molecular modeling (Taylor et al. 2022; Pei et al. 2023). Large molecular language models (LMLMs) (Liu et al. 2023c; Li et al. 2024; Yu et al. 2024a; Park et al. 2024) have recently been explored to harness the instruction-following capabilities of LLMs and shown their effectiveness in diverse applications, such as analyzing molecule structures and drug-related tasks.

Despite these promising efforts, current LMLMs face two key challenges. First, they often produce hallucinated outputs, especially when tasked with complex molecule-to-language generation. Second, they struggle to consistently leverage the complementary strengths of 1D, 2D, and 3D molecular modalities, which encode distinct but interrelated molecular information. For instance, 1D SELFIES represent atom sequences, 2D molecular graphs capture topological connectivity, and 3D conformations provide spatial arrangements critical for reasoning about molecular function. However, existing models typically process these modalities in isolation or shallowly fuse them, limiting their ability to perform molecular reasoning.

To better understand this limitation, we analyze the relationships between molecular tasks and modality contributions. Our empirical study reveals a modality-task dependency, where the usefulness of each modality varies by task, and no single modality universally dominates. These findings motivate the need for a unified framework that can collaboratively integrate all molecular modalities while preserving their structural and spatial nuances.

In this work, we propose CoLLaMo, a large language model-based molecular assistant with a modality-collaborative projector that enables joint reasoning over 1D, 2D, and 3D molecular representations. Specifically, CoLLaMo consists of modality-specific encoders, a large language model, and a modality-collaborative projector that maps representations from different molecular modalities into shared token space through a cross-attention module. To reflect the fine-grained and coherent information of 2D and 3D modalities in the modality-collaborative projector, we present a relation-aware modality-collaborative attention module, which guides the attention map based on 2D structural and 3D structural proximities.

We also present molecule-centric evaluation measurements to address the limitation of evaluation metrics of

*Work was done at Korea University.

[†]Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

LMLMs. Most prior works rely on generic token-level metrics (e.g., BLEU), which are molecule-agnostic and fail to assess chemical correctness or semantic plausibility. To this end, we propose a molecule-centric evaluation framework, including (i) a hallucination assessment metric that quantifies factual inconsistencies in model outputs and (ii) a GPT-based caption quality evaluator that assesses molecular descriptions considering molecular information.

Our extensive experiments show that our CoLLaMo consistently outperforms both LMLMs and LLMs, including GPT-based models such as GPT-4o, o1-mini across a wide range of tasks such as molecule captioning, hallucination assessment, IUPAC naming, motif counting, and molecular QA tasks. Notably, we demonstrate that CoLLaMo maintains strong performance even when certain molecular modalities are missing during inference, showing robustness. Furthermore, our qualitative analysis highlights that CoLLaMo effectively collaborates representations from diverse molecular modalities and well addresses the hallucination by leveraging 1D, 2D, and 3D information, respectively.

Our contributions are summarized as

- We propose a modality-collaborative projector equipped with relation-aware modality-collaborative attention, which facilitates relation-guided information exchange by integrating 2D structural and 3D spatial relations between atoms.
- We introduce CoLLaMo, a large molecular language model based on a modality-collaborative projector, which integrates molecular multimodal representations (1D SELFIES, 2D graphs, and 3D conformations) into unified molecule tokens to fully leverage diverse aspects of molecule information, leading to significantly improved robustness to the molecule hallucination.
- We present a molecule-centric evaluation framework for LMLMs, including an automatic hallucination assessment metric and a GPT-based caption quality evaluator, addressing the limitations of conventional token-based metrics (e.g., BLEU).
- Our experimental results show that CoLLaMo outperforms the strong baselines including GPT-based models GPT-4, GPT-4o, and o1-mini and large molecular language models across various molecule tasks.

2 Related Works

Multimodal Large Language Models. Multimodal Large Language Models (MLLMs) (Dai et al. 2023a; Liu et al. 2023a; Li et al. 2023a; Yuan et al. 2024; Liu et al. 2024) have demonstrated impressive performance on various tasks. However, since most MLLMs depend on a single vision encoder, they suffer from visual hallucinations (Li et al. 2023b; Tong et al. 2024). To address it, multi-encoder strategies have recently been explored, resulting in improved performance and generalization ability (Lin et al. 2024; Kar et al. 2024; Shi et al. 2025; Fan et al. 2024; Zong et al. 2024; Shen et al. 2024; Park et al. 2025). Inspired by these findings, we study the impact of individual molecular modalities on large molecular language models and design a molecular modality unified large molecular language model.

Large Molecular Language Models. Inspired by the success of CLIP (Radford et al. 2021), several studies (Edwards, Zhai, and Ji 2021; Su et al. 2022; Liu et al. 2023b; Lee et al. 2024) have applied molecule-text contrastive learning scheme. Furthermore, multiple studies have adopted LLMs (Zeng et al. 2022; Edwards et al. 2022; Pei et al. 2023; Fang et al. 2024) to understand molecules using the 1D SMILES or SELFIES. Recent works have integrated the language models with 2D graph encoders (Kipf 2017; Xu et al. 2019) or 3D encoders (Zhou et al. 2023; Zhao et al. 2021; Park et al. 2023) with projectors (Zhao et al. 2023; Liu et al. 2023c; Cao et al. 2024; Park et al. 2024). MolCA (Liu et al. 2023c), InstructMol (Cao et al. 2025), and LLaMo (Park et al. 2024) encode molecules with a 2D molecular graph encoder to capture the graph structural relations, while 3D-MoLM (Li et al. 2024) uses 3D encoders to leverage the Euclidean distance information. Despite the individual characteristics of each molecular modality, most works focus on dealing with specific molecular modalities. CoLLaMo aims to leverage 1D, 2D, and 3D molecular information, enhancing generalizability across various tasks.

3 Preliminaries

Large Molecular Language Models. The goal of large molecular language models (LMLMs) is to enable large language models (LLMs) to generate molecular instruction-following responses by incorporating molecular information into LLMs. Formally, the input of LMLMs generally consists of two distinct types of tokens: molecule tokens \mathbf{M} and instruction (text) tokens \mathbf{T} . Given these input tokens, the language model predicts the response $\mathbf{Y} = \{y_i\}_{i=1}^K$ as:

$$p(\mathbf{Y}|\mathbf{M}, \mathbf{T}) = \prod_{i=1}^K p(y_i|\mathbf{M}, \mathbf{T}, \mathbf{Y}_{<i}), \quad (1)$$

where $\mathbf{Y}_{<i}$ denotes a generated token sequence until i -th token and K is the number of generated tokens.

LMLMs generally consist of three modules: (i) molecular encoder, (ii) projector, and (iii) large language model (LLM). The molecular encoder learns representations of input molecules (i.e., SELFIES, molecular graphs, molecular conformations). Then, the projector transfers the molecular representations to molecular tokens comprehensible by the LLM. Finally, the LLM generates the output response given the input molecular and instruction tokens.

Molecular modalities. Molecules can be expressed with 1D molecular strings, 2D molecular graphs, and 3D molecular conformations, each offering distinct information. Given a molecule, m , the 1D molecular string \mathcal{S}_m is typically encoded by molecular descriptors such as SMILES (Weininger 1988) and SELFIES (Krenn et al. 2022), which provide a compact description of atoms and their connectivity in the string form. The 2D molecular graph $\mathcal{G}_m = (\mathcal{V}_m, \mathcal{E}_m)$ represents the molecule’s topological structure, where atoms \mathcal{V}_x are nodes and bonds \mathcal{E}_x are edges. The 3D molecular conformation $\mathcal{C}_m = (\mathcal{V}_m, \mathcal{P}_m)$ assigns spatial coordinates \mathcal{P}_m to each atom, capturing stereochemical and dynamic properties critical for understanding molecules.

Molecular Modality	Captioning (\uparrow)	Descriptive QA (\uparrow)	LogP (\downarrow)	TPSA (\downarrow)	H-L Gap (eV) (\downarrow)	SCF (10^4 eV) (\downarrow)
1D String	25.40	40.87	1.31	38.43	0.46	0.56
2D Graph	32.14	42.74	0.87	20.13	0.31	0.61
3D Conformation	30.72	42.48	0.93	16.98	0.28	0.11

Table 1: Performance comparison of large molecular language models across three molecular modalities.

4 Impact of Molecular Modalities

Each molecular modality captures distinct facets of molecular structure that may influence the performance of molecule-related tasks. To assess the impact of each modality on diverse tasks, we compare model performance using 1D, 2D, and 3D molecular representations. We construct molecule tokens M differently according to the molecular modality and then feed them with instruction tokens T into the LLM to generate a response to the given tasks.

LMLM architecture with 1D molecular modality. To use 1D molecule representation (*i.e.*, Mol-Instructions (Fang et al. 2024)), the molecule tokens M_{1d} are defined as:

$$M_{1d} = \text{Token-Embed}(S_m), \quad (2)$$

where $\text{Token-Embed}(\cdot)$ denotes a token embedding function. In our experiments, we adopt SELFIES as a molecular descriptor and employ a separate vocabulary based on the SELFIES alphabet instead of using the standard BPE tokenizer following existing works (Pei et al. 2023).

LMLM architecture with 2D/3D molecular modalities. While 1D molecular sequences are straightforward to embed using token embeddings, integrating 2D/3D molecular information into LLMs requires additional processing. To address this, recent works (*e.g.*, MolCA (Liu et al. 2023c), LLaMo (Park et al. 2024), 3D-MoLM (Li et al. 2024)) bridge the gap between 2D/3D molecules and natural language by adopting a learnable projector that converts 2D/3D molecular representations into token sequences:

$$M_{2d} = \text{Proj}(g_{2d}(\mathcal{G}_m)), \quad M_{3d} = \text{Proj}(g_{3d}(\mathcal{C}_m))$$

where g_{2d}, g_{3d} are 2D molecular graph and 3D molecular conformer encoder, respectively. $\text{Proj}(\cdot)$ is a projector that aligns the 2D/3D encoder with the language models, enabling them to understand 2D/3D molecules. We use pre-trained GIN (Xu et al. 2019) for the 2D encoder and pre-trained UniMol (Zhou et al. 2023) for the 3D encoder, following previous works. For the projector, we use a MLP similar to existing works (Liu et al. 2023a; Cao et al. 2025).

Evaluation tasks. We evaluate LMLMs across various tasks, including molecule captioning, descriptive property QA, and computed property QA. For the computed property QA, we assess four key properties: LogP, TPSA, HOMO-LUMO gap, and SCF energy. We compare LMLMs that utilize different molecular modality-based tokens (1D, 2D, and 3D). For each task, we compare models that utilize 1D, 2D, or 3D molecular modalities to analyze how modality-specific representations influence performance.

Results. We report the experimental results in Table 1. We observe several important trends. 3D modalities yield the best performance on coordinate-sensitive tasks, such as

HOMO-LUMO gap and SCF energy prediction. This highlights the critical role of spatial information in tasks requiring accurate modeling of electronic structure. 2D graphs perform best on topology-driven tasks, such as LogP prediction, where subgraph patterns and bonding relationships are most informative. 1D SELFIES-based models generally underperform in comparison, suggesting that while sequence representations offer simplicity, they lack the fine-grained structural information required for more complex reasoning. These results confirm that each modality contributes unique information, which is not fully captured by other modalities. This underscores the importance of designing LMLMs that can flexibly integrate and reason over all three modalities to support a wide range of molecular tasks.

5 CoLLaMo

Inspired by the observations in Section 4, we propose CoLLaMo, Collaborative Large Language model-based Molecular assistant that unifies the strength of different molecular modalities with a multi-level molecular modality-collaborative projector. CoLLaMo consists of molecule encoders, the molecular modality-collaborative projector, and an LLM, which is depicted in Figure 1.

5.1 Molecular modality-collaborative projector

The goal of molecular modality-collaborative projector (Co-Proj.) is to leverage 1D SELFIES, 2D graphs, and 3D conformations of the input molecule by encoding them under the shared molecule token space. The Co-Proj. transforms the 1D, 2D, and 3D molecular representations, $Z_{1d} = \text{Token-Embed}(S_x) \in \mathbb{R}^{s \times d_1}$, $Z_{2d} = g_{2d}(\mathcal{G}_x) \in \mathbb{R}^{n \times d_2}$, $Z_{3d} = g_{3d}(\mathcal{C}_x) \in \mathbb{R}^{n \times d_3}$, into a fixed length of molecule token sequence M_{unify} to enable the LLM to understand the molecule. To exploit hierarchical modality information, we adopt a multi-level fusion approach inspired by Park et al. (2024). Specifically, we first process the molecular representations with the attention module as $\hat{Z}_m^{(l)} = \text{Attn}(Z_m^{(l)}, Z_m^{(l)}, Z_m^{(l)}) + Z_m^{(l)}$, where $Z_m^{(l)}$ indicates hidden representations from l -th layer of m -D encoder and $\text{Attn}(Q, K, V)$ is the attention with query Q , key K , and value V . For the 1D representation, we directly use the token embedding output for all layers.

Relation-aware modality-collaborative attention. For 2D and 3D modalities, we adopt modality-collaborative attention instead of standard attention. The 2D graph and 3D conformation have the same atoms but different distinct relational structures (Yu et al. 2024b). To effectively integrate both 2D and 3D relational information, we design a collaborative attention (Co-Attn) that captures their different relational information with the attention bias instead of the standard attention, which is formulated as:

$$\text{Co-Attn}(Q, K, V) = \text{softmax} \left(\frac{(QW_Q)(KW_K)^\top}{\sqrt{d}} + \Phi_{2d} + \Phi_{3d} \right) VW_V, \quad (3)$$

where Φ_{2d}, Φ_{3d} are attention bias based on 2D and 3D structural information, respectively.

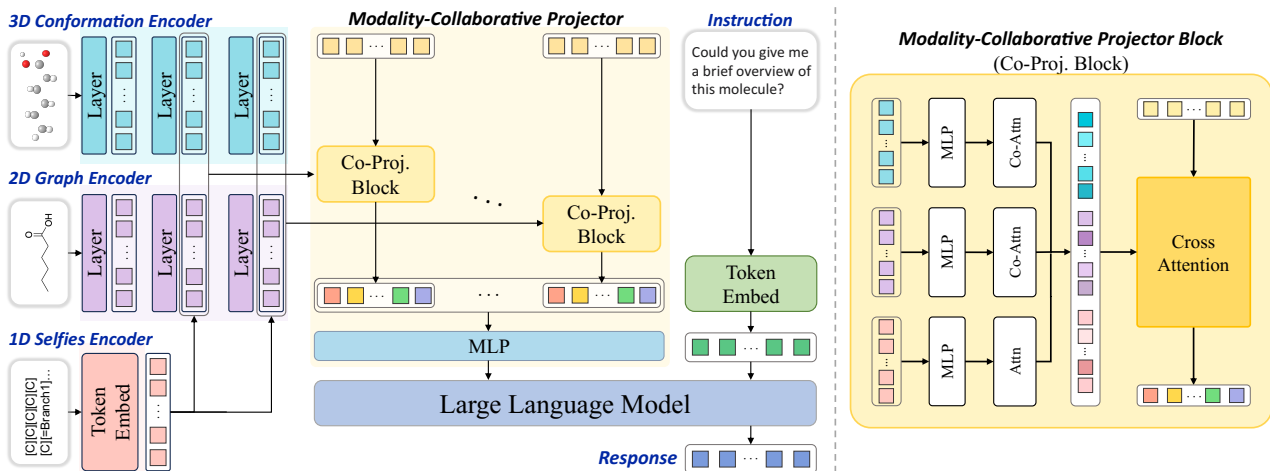


Figure 1: Overall architecture of CoLLaMo. Our framework, CoLLaMo, consists of encoders for 1D SELFIES, 2D graphs, and 3D conformations, a molecular modality-collaborative projector, and a large language model. Our framework first encodes 1D, 2D, and 3D molecular representations and then converts the encoded outputs into unified molecule tokens using the modality-collaborative projector. Finally, the large language model generates the response given the input molecule and instruction.

Specifically, we define Φ_{2d} using the shortest path distance between atom pairs, a standard measure for capturing graph connectivity. The shortest path distance-based attention bias Φ_{2d} is defined as a learnable scalar value, where the value is determined by the shortest path between atoms for each pair of atoms. For 3D structural information, we use Euclidean distance to leverage the spatial information in 3D space. Following Yu et al. (2024b), we first use the Gaussian Basis Kernel function (Scholkopf et al. 1997) to process the distance, $\psi_{(i,j)}^k = -\frac{1}{\sqrt{2\pi}|\sigma^k|} \exp\left(-\frac{1}{2}\left(\frac{\|r_i - r_j\| - \mu^k}{|\sigma^k|}\right)^2\right)$, where $\mu^k, \sigma^k \in \mathbb{R}$ are learnable parameters. Then, the 3D distance-based attention bias $\Phi_{3d} \in \mathbb{R}^{n \times n}$ is calculated as $\Phi_{3d} = \text{MLP}\left([\Psi^1; \dots; \Psi^K]\right)$, where Ψ^k indicates the processed distance matrix based on the k -th Gaussian kernel and K is the number of Gaussian Basis kernel.

Modality token unification via cross-attention. To produce a condensed and unified molecule tokens, we adopt cross-attention with the learnable tokens $P^{(l)} = [p_1^{(l)}; \dots; p_b^{(l)}] \in \mathbb{R}^{d \times b}$ for $l = 1, \dots, L$ to deal with an arbitrary number of atoms. Specifically, the representations $\hat{Z}_{1d}^{(l)}, \hat{Z}_{2d}^{(l)}, \hat{Z}_{3d}^{(l)}$ are unified into a fixed number of tokens $\hat{P}^{(l)}$ as:

$$\hat{P}^{(l)} = \text{Attn}\left(P^{(l)}, \hat{Z}_{\text{concat}}^{(l)}, \hat{Z}_{\text{concat}}^{(l)}\right),$$

where $\hat{Z}_{\text{concat}}^{(l)}$ is defined as $[\hat{Z}_{1d}^{(l)}; \hat{Z}_{2d}^{(l)}; \hat{Z}_{3d}^{(l)}]$. We then construct unified molecule tokens $M_{\text{unify}} \in \mathbb{R}^{d \times (b \cdot L)}$ by feeding tokens $\hat{P}^{(l)}$ into MLP as: $M_{\text{unify}} = \text{MLP}\left([\hat{P}^{(1)}; \dots; \hat{P}^{(L)}]\right)$. This enables the molecule tokens to unify different molecular modalities in a shared semantic space, which promotes modality collaboration.

Molecular modality dropout & embedding. To prevent over-reliance on a specific modality, we apply modality dropout during pretraining by randomly masking each modality’s representation. This encourages the model to learn more robust representations by leveraging the complementary information provided by the remaining modalities. To further inform the model of the active modalities, we incorporate modality-specific learnable embeddings: $P^{(l)} = P_{\text{base}}^{(l)} + \sum_{m \in \text{active}} P_m^{(l)}$, where $\text{active} \subseteq \{1d, 2d, 3d\}$. For example, when the 2D molecular modality is dropped, we use the learnable query tokens $P^{(l)} = P_{\text{base}}^{(l)} + P_{1d}^{(l)} + P_{3d}^{(l)}$. By jointly leveraging both modality dropout and modality embeddings, our approach enhances the model’s generalization ability across tasks without being overly dependent on specific molecular modality.

5.2 Training CoLLaMo

Following most existing works, our CoLLaMo is trained in multiple stages.

Stage 1. Pretraining for molecule-language alignment.

The goal of the first stage is to align molecule encoders with a large language model. We train our Co-Proj. and molecule encoders while the LLM is being frozen. In this stage, we use two types of molecule-caption pair datasets: a molecule-caption pair dataset from Fang et al. (2024) and the GPT-enriched data presented in Li et al. (2024).

Stage 2. Molecule-language instruction-tuning.

The second stage aims to improve the molecule-language instruction-following capabilities of LLMs and enable them to understand molecules. We train both the Co-Proj. and LLM equipped with LoRA (Hu et al. 2022) while keeping the molecule encoders frozen. In this stage, we employ a diverse set of datasets with various instructions: training splits of molecule captioning, descriptive property QA, computed property QA, motif counting, and IUPAC prediction.

Model	LLM	Molecule Captioning		IUPAC Prediction		Motif Counting	
		BLEU (\uparrow)	METEOR (\uparrow)	BLEU (\uparrow)	METEOR (\uparrow)	MAE (\downarrow)	Validity (\uparrow)
GPT-4	GPT-4	0.8	16.7	28.3	50.4	5.78	100%
GPT-4 (ICL)	GPT-4	27.0	52.2	51.7	62.7	2.51	100%
GPT-4o	GPT-4o	1.6	15.6	11.2	37.2	4.21	95%
GPT-4o (ICL)	GPT-4o	31.1	57.1	39.4	55.7	2.45	100%
o1-mini	o1-mini	1.8	16.6	2.4	26.5	9.62	100%
o1-mini (ICL)	o1-mini	27.2	53.0	29.4	50.5	2.72	100%
LLaMA2	LLaMA2-7B	0.0	14.1	0.0	0.4	N/A*	-
Mol-Instructions	LLaMA2-7B	14.3	25.4	-	-	-	-
LLaMo	LLaMA2-7B	38.9	67.1	56.0	73.2	-	-
CoLLaMo (Ours)	LLaMA2-7B	44.5	70.5	59.8	77.0	1.31	100%

Table 2: Evaluation results of generalist models on molecule captioning, hallucination, IUPAC prediction, and motif counting. ICL denotes in-context learning. * The result is not available since LLaMA2 fails generating numerical outputs.

Model	Descriptive Property QA		Computed Property QA											
	BLEU	METEOR	LogP		TPSA (\AA^2)		HOMO (eV)		LUMO (eV)		H-L Gap (eV)		SCF (10^4 eV)	
			MAE	Validity	MAE	Validity	MAE	Validity	MAE	Validity	MAE	Validity	MAE	Validity
Uni-Mol (Non-LM)	-	-	0.59	100%	13.48	100%	0.32	100%	0.35	100%	0.21	100%	0.45	100%
<i>Specialist</i>														
Llama2-7B	23.24	46.87	1.45	95%	15.87	92%	1.24	96%	1.04	95%	0.88	92%	0.70	99%
2D-MoLM	25.09	50.92	0.88	96%	13.52	92%	0.92	98%	0.80	96%	0.67	93%	0.71	99%
3D-MoLM \dagger	24.47	51.33	0.95	96%	10.26	94%	0.45	98%	0.36	96%	0.41	94%	0.39	99%
3D-MoLM	26.13	52.15	0.66	97%	9.71	93%	0.26	97%	0.25	94%	0.28	94%	0.35	99%
CoLLaMo (Ours)	45.62	60.06	0.37	100%	4.93	100%	0.08	100%	0.08	100%	0.09	100%	0.02	100%
<i>Generalist</i>														
Llama2-7B*	21.16	43.17	2.10	85%	27.11	76%	2.87	70%	1.89	71%	1.86	70%	3.84	23%
Llama2-7B	22.81	46.39	1.78	93%	17.07	90%	1.89	90%	1.26	90%	1.25	91%	0.87	99%
2D-MoLM	24.57	50.08	1.36	94%	12.47	89%	1.52	93%	1.13	92%	1.09	88%	0.96	99%
3D-MoLM \dagger	24.44	50.81	0.92	92%	11.14	92%	0.65	94%	0.41	92%	0.55	89%	0.49	99%
3D-MoLM	26.08	51.93	0.78	95%	10.90	90%	0.35	95%	0.36	93%	0.32	90%	0.38	98%
CoLLaMo (Ours)	45.62	59.89	0.63	100%	7.79	100%	0.13	100%	0.14	100%	0.13	100%	0.09	100%

Table 3: Experimental results for descriptive property QA and computed property QA. For computed property QA, MAE is used for the evaluation metric. The validity is also reported since LMs sometimes fail to generate valid numerical responses. * represents the model without fine-tuning. \dagger denotes the variant of 3D-MoLM that is pretrained using original PubChem text without GPT-3.5 enrichment.

6 Advancing Molecule-Aware Evaluation

Hallucination assessment. Although recent large molecule-language models (LMLMs) (Park et al. 2024; Liu et al. 2023c; Li et al. 2024) have shown progress in generating molecule-related descriptions, their outputs often contain factual inaccuracies or spurious content. Despite the importance of the hallucination of LMLMs, previous studies have underexplored the evaluation metrics for the hallucination. So, we introduce two quantitative metrics specifically designed to evaluate the factual consistency of LMLM outputs.

Inspired by CHAIR metric (Rohrbach et al. 2018), we propose **CHARM** (Caption Hallucination Assessment with Relevance to Molecule), a metric designed to assess hallucination in molecule-language models. CHARM quantifies the proportion of mentioned molecular entities that are factually incorrect or not grounded in the input molecule. Specifically, we first extract molecular entities by applying BERN2 (Sung et al. 2022), a widely-used and effective Named Entity Recognition (NER) system. To ensure

the reliability of the extracted entities, we filtered out low-confidence predictions and only retained entities with a confidence score of 0.9 or higher. With extracted entities, we calculate metrics as follows:

$$\text{CHARM} = \frac{|\{\text{hallucinated entities}\}|}{|\{\text{all mentioned entities}\}|}, \quad (4)$$

This value represents the proportion of hallucinated entities among all entities mentioned. As CHARM corresponds to a precision-oriented view of factuality, we further propose **RCHARM**, a complementary, recall-oriented metric that captures the proportion of relevant entities from the ground truth that are omitted in the generated output:

$$\text{RCHARM} = \frac{|\{\text{not mentioned g.t. entities}\}|}{|\{\text{all g.t. entities}\}|}. \quad (5)$$

Together, CHARM and RCHARM enable a more comprehensive and balanced assessment of hallucination in molecule-language models.

Model	Molecular Entity Hallucination		LLM Score
	CHARM (↓)	RCHARM (↓)	Avg. Score (↑)
GPT-4	98.4	94.3	1.95
GPT-4 (ICL)	77.1	76.9	2.04
GPT-4o	99.0	98.4	1.99
GPT-4o (ICL)	74.0	75.1	2.15
o1-mini	99.1	98.8	2.17
o1-mini (ICL)	83.5	82.7	2.01
LLaMo	64.7	67.2	2.17
CoLLaMo (Ours)	58.5	59.9	2.52

Table 4: Evaluation with our evaluation protocols, such as CHARM/RCHARM and GPT-based LLM score. GPT-4o is used as a judge.

LLM-based molecule description assessment. Traditional evaluation metrics, such as BLEU and ROUGE, are occasionally inadequate for assessing molecular descriptions, where ground-truth references may be underspecified or lack molecular precision. To overcome these limitations, we adopt a large language model (LLM)-based evaluation approach, following recent trends in LLM-as-a-judge setups (Zheng et al. 2023). Specifically, we employ GPT-4o to assess the quality of generated molecular descriptions from CoLLaMo and baseline models. The evaluation prompt is structured to include: (i) the input molecule in SELFIES format, (ii) the ground-truth description, and (iii) the model-generated output. The LLM judge is asked to score the response based on two criteria: factual informativeness and alignment with the ground truth. Based on criteria, the model assigns a human-likeness score on a 0–5 scale.

7 Experiments

7.1 Experimental Settings

Benchmarks. To validate the effectiveness of our CoLLaMo, we compare the performance of baseline models on diverse tasks such as **1)** molecule captioning, **2)** molecule hallucination assessment, **3)** descriptive property QA, **4)** computed property QA, **5)** IUPAC prediction, and **6)** motif counting.

Implementation details. We utilize the LLaMA 2 Chat 7B model as our base language model. For the analysis, including the ablation studies, we employ a **short training schedule** consisting of 64k iterations of pre-training followed by 150k iterations of instruction tuning. For the final models, we adopt an **extended training schedule** with 64k iterations of pre-training and 900k iterations of instruction tuning. We conduct experiments on two variants of CoLLaMo: generalist, which is a single model trained on all the tasks, and specialist, which is trained for each task following other works (Li et al. 2024).

7.2 Experimental Results

Molecule captioning. We evaluate CoLLaMo and baselines on molecule captioning in Table 2. The table shows the superiority of CoLLaMo on molecule captioning. In comparison to GPT-based models such as GPT-4, GPT-4o, and o1-mini, our CoLLaMo achieves the best performance.

Mol. modality	Proj.	Captioning	Motif Count	Desc. QA	H-L Gap		
						1D	2D
✓		33.2	1.82	42.9	0.31		
	✓	38.0	1.41	43.9	0.27		
		35.8	1.65	42.6	0.23		
✓	✓	35.7	1.59	43.6	0.21		
✓	✓	✓	✓	40.1	1.43	44.8	0.18

Table 5: Ablation study on molecular modality. C.P denotes molecular modality-collaborative projector.

Component	Captioning	Motif Count	Desc. QA	H-L Gap
	BLEU (↑)	MAE (↓)	BLEU (↑)	MAE (↓)
w/o Co-Attention	36.6	1.44	43.9	0.19
w/o Modality embedding	38.9	1.51	44.3	0.19
w/o Modality dropout	37.8	1.49	44.5	0.18
CoLLaMo (Ours)	40.1	1.43	44.8	0.18

Table 6: Ablation study on components of CoLLaMo.

IUPAC prediction and motif counting. We provide additional experimental results on IUPAC name prediction and motif counting tasks in Table 2. Our CoLLaMo consistently outperforms GPT-based models across both tasks with its effective molecular understanding capabilities.

Molecular property QA. Table 3 presents the experimental results on open-text descriptive and computed property QA tasks. From the table, our CoLLaMo outperforms the baselines using 2D molecular graphs (2D-MoLM) or 3D conformations (3D-MoLM) in both *Generalist* and *Specialist* settings. In particular, our CoLLaMo achieves a significant performance improvement of 19.54 in BLEU and 7.96 in METEOR compared to 3D-MoLM on the descriptive property QA task in the *Generalist* setting. These results highlight that 1D, 2D, and 3D modalities have effectively collaborated. Additionally, CoLLaMo surpasses molecule-specialized non-LM model Uni-Mol on all properties in the *Specialist* setting, underscoring its ability to combine the contextual knowledge from diverse molecular modalities.

Molecular Entity hallucination. We evaluate the extent of hallucination using our proposed metrics, CHARM and RCHARM in Table 4. The table shows that CoLLaMo outperforms GPT-based models and LLaMo on both CHARM and RCHARM metrics, which demonstrates that CoLLaMo successfully collaborates with diverse molecular modalities and complements information with each other.

LLM-based evaluation. To further validate the quality of generated outputs, we employed an LLM-as-a-judge setup using GPT-4o to evaluate responses from CoLLaMo and baseline models. The average human-likeness scores (scale: 0–5) are summarized in Table 4. From the table, CoLLaMo achieves the highest score, demonstrating its superior molecule understanding and generation quality from both automatic and LLM-based evaluations.

7.3 Analysis

Impact of molecular modalities and the molecular modality-collaborative projector. To assess the contribution of different molecular modalities, we evaluate the model using 1D only, 2D only, 3D only, and unification of 1D, 2D,

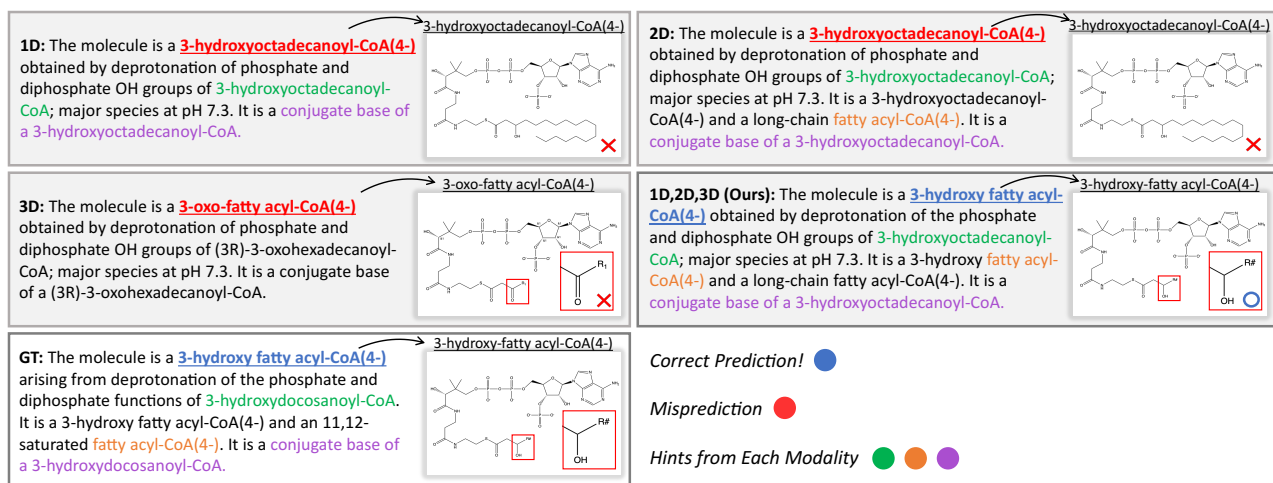


Figure 2: Examples of molecule captioning given the molecule (CHEBI:65102). The molecule structures of underlined texts in generated descriptions are illustrated in the white box.

Co-Proj.	Adopted Mol. modality	Captioning BLEU	Motif Count MAE	Desc. QA BLEU	H-L Gap MAE
w/o Co-Proj.	1D	2.4	N/A	12.8	N/A
	1D+2D	3.3	N/A	13.8	N/A
	1D+2D+3D	35.7	1.59	43.6	0.21
w/ Co-Proj. (Ours)	1D	28.0	1.94	42.4	0.31
	1D+2D	39.1	1.43	44.0	0.24
	1D+2D+3D	40.1	1.43	44.8	0.18

Table 7: Experimental results under **missing modality settings**. Adopted Mol. modality denotes molecular modality employed during the inference phase. Co-Proj. denotes molecular modality-collaborative projector.

and 3D molecular modalities with and without the molecular modality-collaborative projector (C.P.) in Table 5. For the model without C.P., the cross-attention is separately applied for each molecular modality. The results demonstrate that integrating 1D, 2D, and 3D modalities through the modality-collaborative projector consistently improves the performance across all tasks except motif counting.

Effectiveness of components in CoLLaMo. Table 6 shows the results of ablation studies on the key components of CoLLaMo, including Co-Attention, modality dropout, and modality embedding. The experimental results reveal that all components contribute to the performance improvement of our CoLLaMo. Notably, when comparing CoLLaMo with and without Co-Attention, we observe a significant performance improvement of 3.5 in the BLEU score on the molecule captioning task. Additionally, the inclusion of Co-Attention yields substantial gains in performance, highlighting the importance of this regularization technique.

Robustness under missing molecular modality. We further evaluate CoLLaMo’s robustness in scenarios with incomplete molecular modalities by comparing the model with and without a molecular collaborative-modality projector (Co-Proj.), as presented in Table 7. The robustness is assessed by utilizing an incomplete set of molecular modalities during the inference phase. We note that all molecular

modalities are used in the training phase. The results reveal that even when only a single modality is available during inference, CoLLaMo maintains strong performance, demonstrating its robustness, while the model without Co-Projector is significantly degraded. Since CoLLaMo embeds different molecular modalities into shared representations, it preserves molecule information given the incomplete molecular modalities. These findings highlight CoLLaMo’s capacity to generalize well in real-world scenarios, where modality information may be incomplete or unavailable.

Qualitative analysis. Figure 2 presents a comparison between the ground truth (GT) caption and the captions generated by models trained with individual molecular modalities (1D string, 2D graph, 3D conformation) and all modalities combined. As shown in the figure, CoLLaMo generates the most accurate molecular description, leveraging all molecular modalities. The GT caption describes the molecule with “3-hydroxy fatty acyl-CoA(4-)”. While the model with a single molecular modality wrongly explains the molecule, CoLLaMo with all modalities generates an accurate description with “3-hydroxy fatty acyl-CoA(4-)” by successfully integrating the information from each modality.

8 Conclusion

We propose CoLLaMo, a molecular modality **C**ollaborative projector-based **L**arge **L**anguage model-based **M**olecular assistant, which integrates 1D, 2D, and 3D modalities under a unified molecule token space. We also present new measurements to assess the hallucination in the outputs generated by the language models, addressing the limitations of existing measurements. Our experiments demonstrate the effectiveness of CoLLaMo, outperforming GPT-based models such as GPT-4, GPT-4o, and o1-mini on diverse tasks, including molecule captioning, IUPAC prediction, and motif counting. In particular, our CoLLaMo shows good experimental results under missing modality settings, highlighting that it effectively integrates 1D, 2D, and 3D modalities.

Acknowledgements

This research was supported by the ASTRA Project through the National Research Foundation (NRF) funded by the Ministry of Science and ICT (No. RS-2024-00439619).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv:2303.08774*.
- Cao, H.; Liu, Z.; Lu, X.; Yao, Y.; and Li, Y. 2025. Instruct-mol: Multi-modal integration for building a versatile and reliable molecular assistant in drug discovery. In *COLING*.
- Cao, H.; Shao, Y.; Liu, Z.; Tang, X.; Yao, Y.; and Li, Y. 2024. PRESTO: Progressive Pretraining Enhances Synthetic Chemistry Outcomes. In *EMNLP-Findings*.
- Chen, S.-A.; Miculicich, L.; Eisenschlos, J.; Wang, Z.; Wang, Z.; Chen, Y.; Fujii, Y.; Lin, H.-T.; Lee, C.-Y.; and Pfister, T. 2024. Tablerag: Million-token table understanding with language models. In *NeurIPS*.
- Chu, J.; Lee, S.; and Kim, H. J. 2025. PRESTO: Preimage-Informed Instruction Optimization for Prompting Black-Box LLMs. In *NeurIPS*.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P.; and Hoi, S. C. H. 2023a. InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning. In *NeurIPS*.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P. N.; and Hoi, S. 2023b. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *NeurIPS*.
- Edwards, C.; Lai, T.; Ros, K.; Honke, G.; Cho, K.; and Ji, H. 2022. Translation between molecules and natural language. In *EMNLP*.
- Edwards, C.; Zhai, C.; and Ji, H. 2021. Text2mol: Cross-modal molecule retrieval with natural language queries. In *EMNLP*.
- Fan, X.; Ji, T.; Jiang, C.; Li, S.; Jin, S.; Song, S.; Wang, J.; Hong, B.; Chen, L.; Zheng, G.; et al. 2024. MouSi: Poly-Visual-Expert Vision-Language Models. In *COLM*.
- Fang, Y.; Liang, X.; Zhang, N.; Liu, K.; Huang, R.; Chen, Z.; Fan, X.; and Chen, H. 2024. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. In *ICLR*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. Lora: Low-rank adaptation of large language models. In *ICLR*.
- Kar, O. F.; Tonioni, A.; Poklukar, P.; Kulshrestha, A.; Zamir, A.; and Tombari, F. 2024. BRAVE: Broadening the visual encoding of vision-language models. In *ECCV*.
- Kim, J.; Bae, M.; Lee, S.; Yoon, J.; and Kim, H. J. 2026. TabFlash: Efficient Table Understanding with Progressive Question Conditioning and Token Focusing. In *AAAI*.
- Kipf, T. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Krenn, M.; Ai, Q.; Barthel, S.; Carson, N.; Frei, A.; Frey, N. C.; Friederich, P.; Gaudin, T.; Gayle, A. A.; Jablonka, K. M.; et al. 2022. SELFIES and the future of molecular string representations. *Patterns*, 3(10).
- Lee, N.; Laghuvarapu, S.; Park, C.; and Sun, J. 2024. Vision Language Model is NOT All You Need: Augmentation Strategies for Molecule Language Models. In *CIKM*.
- Li, L.; Yin, Y.; Li, S.; Chen, L.; Wang, P.; Ren, S.; Li, M.; Yang, Y.; Xu, J.; Sun, X.; et al. 2023a. M³ IT: A Large-Scale Dataset towards Multi-Modal Multilingual Instruction Tuning. *arXiv:2306.04387*.
- Li, S.; Liu, Z.; Luo, Y.; Wang, X.; He, X.; Kawaguchi, K.; Chua, T.-S.; and Tian, Q. 2024. Towards 3d molecule-text interpretation in language models. In *ICLR*.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J.-R. 2023b. Evaluating object hallucination in large vision-language models. In *EMNLP*.
- Lin, Z.; Liu, C.; Zhang, R.; Gao, P.; Qiu, L.; Xiao, H.; Qiu, H.; Lin, C.; Shao, W.; Chen, K.; et al. 2024. Sphinx: SPHINX: A Mixer of Weights, Visual Embeddings and Image Scales for Multi-modal Large Language Models. In *ECCV*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024. Improved baselines with visual instruction tuning. In *CVPR*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual instruction tuning. In *NeurIPS*.
- Liu, S.; Nie, W.; Wang, C.; Lu, J.; Qiao, Z.; Liu, L.; Tang, J.; Xiao, C.; and Anandkumar, A. 2023b. Multi-modal molecule structure-text model for text-based retrieval and editing. *Nat. Mach. Intell.*, 5(12): 1447-1457.
- Liu, Z.; Li, S.; Luo, Y.; Fei, H.; Cao, Y.; Kawaguchi, K.; Wang, X.; and Chua, T.-S. 2023c. MolCA: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In *EMNLP*.
- Park, J.; Bae, M.; Ko, D.; and Kim, H. J. 2024. Llamol: large language model-based molecular graph assistant. In *NeurIPS*.
- Park, J.; Lee, S.; Kim, S.; Xiong, Y.; and Kim, H. J. 2023. Self-positioning point-based transformer for point cloud understanding. In *CVPR*.
- Park, J.; Na, J.; Kim, J.; and Kim, H. J. 2025. DeepVideoR1: Video Reinforcement Fine-Tuning via Difficulty-aware Regressive GRPO. In *NeurIPS*.
- Pei, Q.; Zhang, W.; Zhu, J.; Wu, K.; Gao, K.; Wu, L.; Xia, Y.; and Yan, R. 2023. Biot5: Enriching cross-modal integration in biology with chemical knowledge and natural language associations. In *EMNLP*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Rohrbach, A.; Hendricks, L. A.; Burns, K.; Darrell, T.; and Saenko, K. 2018. Object Hallucination in Image Captioning. In *EMNLP*.

- Scholkopf, B.; Sung, K.-K.; Burges, C. J.; Girosi, F.; Niyogi, P.; Poggio, T.; and Vapnik, V. 1997. Comparing support vector machines with Gaussian kernels to radial basis function classifiers. *TSP*, 45(11): 2758–2765.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv:2402.03300*.
- Shen, L.; Chen, G.; Shao, R.; Guan, W.; and Nie, L. 2024. MoME: Mixture of Multimodal Experts for Generalist Multimodal Large Language Models. In *NeurIPS*.
- Shi, M.; Liu, F.; Wang, S.; Liao, S.; Radhakrishnan, S.; Huang, D.-A.; Yin, H.; Sapra, K.; Yacoob, Y.; Shi, H.; et al. 2025. Eagle: Exploring The Design Space for Multimodal LLMs with Mixture of Encoders. In *ICLR*.
- Su, B.; Du, D.; Yang, Z.; Zhou, Y.; Li, J.; Rao, A.; Sun, H.; Lu, Z.; and Wen, J.-R. 2022. A molecular multimodal foundation model associating molecule graphs with natural language. *arXiv preprint arXiv:2209.05481*.
- Sung, M.; Jeong, M.; Choi, Y.; Kim, D.; Lee, J.; and Kang, J. 2022. BERN2: an advanced neural biomedical named entity recognition and normalization tool. *Bioinformatics*, 38(20): 4837–4839.
- Taylor, R.; Kardas, M.; Cucurull, G.; Scialom, T.; Hartshorn, A.; Saravia, E.; Poulton, A.; Kerkez, V.; and Stojnic, R. 2022. Galactica: A large language model for science. *arXiv:2211.09085*.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv:2312.11805*.
- Tong, S.; Liu, Z.; Zhai, Y.; Ma, Y.; LeCun, Y.; and Xie, S. 2024. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *CVPR*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*.
- Weininger, D. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.*, 28(1): 31–36.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How powerful are graph neural networks? In *ICLR*.
- Yang, A.; Zhang, B.; Hui, B.; Gao, B.; Yu, B.; Li, C.; Liu, D.; Tu, J.; Zhou, J.; Lin, J.; et al. 2024. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv:2409.12122*.
- Yu, B.; Baker, F. N.; Chen, Z.; Ning, X.; and Sun, H. 2024a. Llamol: Advancing large language models for chemistry with a large-scale, comprehensive, high-quality instruction tuning dataset. In *COLM*.
- Yu, Q.; Zhang, Y.; Ni, Y.; Feng, S.; Lan, Y.; Zhou, H.; and Liu, J. 2024b. Multimodal Molecular Pretraining via Modality Blending. In *ICLR*.
- Yuan, Y.; Li, W.; Liu, J.; Tang, D.; Luo, X.; Qin, C.; Zhang, L.; and Zhu, J. 2024. Osprey: Pixel understanding with visual instruction tuning. In *CVPR*.
- Zeng, Z.; Yao, Y.; Liu, Z.; and Sun, M. 2022. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nat. Commun.*, 13(1): 862.
- Zhang, T.; Yue, X.; Li, Y.; and Sun, H. 2024. Tablellama: Towards open large generalist models for tables. In *NAACL*.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021. Point transformer. In *ICCV*.
- Zhao, H.; Liu, S.; Chang, M.; Xu, H.; Fu, J.; Deng, Z.; Kong, L.; and Liu, Q. 2023. Gimlet: A unified graph-text model for instruction-based molecule zero-shot learning. In *NeurIPS*.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *NeurIPS*.
- Zhou, G.; Gao, Z.; Ding, Q.; Zheng, H.; Xu, H.; Wei, Z.; Zhang, L.; and Ke, G. 2023. Uni-mol: A universal 3d molecular representation learning framework. In *ICLR*.
- Zong, Z.; Ma, B.; Shen, D.; Song, G.; Shao, H.; Jiang, D.; Li, H.; and Liu, Y. 2024. Mova: Adapting mixture of vision experts to multimodal context. In *NeurIPS*.