Hierarchical Graph Tokenization for Molecule-Language Alignment

Yongqiang Chen¹²³ Quanming Yao⁴ Juzheng Zhang⁵ James Cheng³ Yatao Bian⁶

Abstract

Recently, there has been a surge of interest in extending the success of large language models (LLMs) from texts to molecules. Most existing approaches adopt a graph neural network to represent a molecule as a series of node tokens for molecule-language alignment, which, however, have overlooked the inherent hierarchical structures in molecules. Notably, higherorder molecular structures contain rich semantics of functional groups, which encode crucial biochemical functionalities of the molecules. We show that neglecting the hierarchical information in tokenization will lead to subpar moleculelanguage alignment and severe hallucination. To address this limitation, we propose HIerarchical GrapH Tokenization (HIGHT). HIGHT employs a hierarchical graph tokenizer that encodes the hierarchy of atom, motif, and molecular levels of informative tokens to improve the molecular perception of LLMs. HIGHT also adopts an augmented instruction tuning dataset, enriched with the hierarchical graph information, to further enhance the molecule-language alignment. Extensive experiments on 14 real-world benchmarks verify the effectiveness of HIGHT in reducing hallucination by 40%, and significant improvements in various molecule-language downstream tasks. The project is available at https: //higraphllm.github.io/.

1. Introduction

Large language models (LLMs) have demonstrated impressive capabilities in understanding and processing natural languages (Radford et al., 2019; OpenAI, 2022; Touvron et al., 2023a; Bubeck et al., 2023). Recently, there has been a surge of interest in extending the capabilities of LLMs to graph-structured data (Jin et al., 2023; Li et al., 2023d; Wei et al., 2024; Mao et al., 2024; Fan et al., 2024), particularly molecular graphs (Zhao et al., 2023; Cao et al., 2023). Inspired by the success of large vision-language models (Zhang et al., 2024b; Liu et al., 2023a), recent efforts in developing large graph-language models (LGLMs) typically adopt a graph neural network (GNN) (Xu et al., 2019) to tokenize molecules as a series of node embeddings (or node tokens), and then leverage an adapter such as a Multi-layer perceptron (MLP) or a Q-former (Li et al., 2023a) to transform the node tokens into those compatible with LLMs (Fan et al., 2024). To bridge the gap between the graph and language modalities, LGLMs will undergo a molecule-language instruction tuning with the molecular graph and the corresponding captions describing the molecules (Jin et al., 2023; Li et al., 2023d; Fan et al., 2024).

Despite recent progress, the tokenization in existing LGLMs neglects the essential hierarchical structures inherent in molecular graphs. In particular, in molecular graphs, the high-order substructures, such as motifs or functional groups, encode rich semantics of the biochemical functionalities of the molecules (Milo et al., 2002; Bohacek et al., 1996; Sterling & Irwin, 2015). For example, the presence of a hydroxide functional group ("-OH") often indicates a higher water solubility. Therefore, such substructural cues are essential for enabling LLMs to reason about the molecules in a chemically meaningful way. However, existing LGLMs mostly tokenize molecules solely at the atom (node) level, and feed LLMs with only node-level tokens. Consequently, it requires LLMs to implicitly infer the underlying substructures during the instruction tuning stage. The absence of the critical substructures not only increases the unnecessary burdern on the LLMs, but also leads to misaligned representations and a higher likelihood of hallucinations in downstream tasks. To quantify the issue, we introduce a diagnostic benchmark, called MotifHallu, which evaluates the perception ability of LGLMs about the existence of common functional groups. Surprisingly, we find that existing LGLMs often produce false-positive predictions (i.e., keep answering "Yes" for any functional groups), highlighting a critical limitation in current graph to-

Most of the works were done when Yongqiang Chen was a PhD student at CUHK. ¹MBZUAI ²Carnegie Mellon University ³The Chinese University of Hong Kong ⁴Tsinghua University ⁵University of Maryland, College Park ⁶Department of Computer Science, National University of Singapore. Correspondence to: Yatao Bian
bianyt@comp.nus.edu.sg>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

Hierarchical Graph Tokenization for Molecule-Language Alignment



Figure 1. (a) **Illustration of HIGHT**: Given a molecule (i.e., PubChem ID 3, *5,6-Dihydroxycyclohexa-1,3-diene-1-carboxylic acid*), HIGHT detects the motifs and incorporates the "supernodes" for each motif (The whole graph is also considered as a "super motif".). Then, HIGHT tokenizes the molecule into both node-level (i.e., atoms) and motif-level (i.e., functional groups) tokens. The hierarchical view enables LLMs to align the molecular structures and the language descriptions of the molecule better. (*b*) **Performance Overview**: HIGHT significantly reduces the hallucination of LGLMs and improves the downstream performance across various molecule-centric tasks. Due to the heterogeneity of the evaluation metrics in each task, we perform some transformations on the numerical values. In MotifHallu, we report the macro F1 scores. For Property Classification and Molecular Caption, we report the averaged scores of all the subtasks or submetrics. For Property Regression, we normalize the values to the range between 1 and 100, i.e., for *a*, the reported number is 0.5/a. For Chemical Reaction Prediction, we report the averaged values of BLEU, RDK, MACCS, and MORGAN.

kenization strategies (Sec. 3.2). This observation motivates the following research question:

Is there a feasible approach to integrate the intrinsic hierarchical molecular information into LLMs?

To tackle the problem, we propose a new molecule-language alignment strategy called **HI**erarchical **GrapH** Tokenization (HIGHT). As illustrated in Fig. 1, HIGHT adopts a hierarchical graph tokenizer and a hierarchical molecular instruction tuning dataset to facilitate a better alignment of molecule and language modalities. Specifically, inspired by the success of hierarchical GNNs in molecular representation learning (Zhang et al., 2021; Zang et al., 2023; Inae et al., 2023; Luong & Singh, 2023), HIGHT transforms the original molecular graph into a hierarchical graph with motif-level and molecule-level nodes added in. Then, HIGHT employs a Vector Quantized-Variational AutoEncoder (VQVAE) to obtain atom-level, motif-level, and molecule-level tokens separately with the self-supervised tasks (Zang et al., 2023).

In addition, to further encourage the encoding and alignment of hierarchical information, HIGHT augments the original molecular instruction tuning dataset with motif-level descriptions. Our contributions can be summarized as follows:

- To the best of our knowledge, we are the *first to incorporate the hierarchical graph information* into LGLMs, with the consideration of both the architecture-level and the instruction tuning data.
- To facilitate the molecule-language alignment study,

we also propose the *first hallucination benchmark* MotifHallu, synthesized through question-answering based on common functional groups.

• We conduct extensive experiments with 14 real-world benchmarks. The results show that HIGHT significantly reduces the hallucination on MotifHallu by up to 40% and consistently improves the performances on downstream molecule-language tasks.

Hence, HIGHT together with MotifHallu and HiPubChem, lay the solid foundation for developing graph foundation models via graph-language alignment.

2. Preliminaries

Large Graph-Language Models. As LLMs have demonstrated great capabilities across a wide range of natural language tasks, there has been an increasing interest in extending LLMs to broader applications where the text data are associated with the structure information (i.e., graphs) (Jin et al., 2023; Li et al., 2023d; Wei et al., 2024; Mao et al., 2024; Fan et al., 2024). A graph can be denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with a set of *n* nodes $v \in \mathcal{V}$ and a set of *m* edges $(u, v) \in \mathcal{E}$. Each node *u* has node attributes as $\boldsymbol{x}_u \in \mathbb{R}^d$ and each edge (u, v) has edge attributes $e_{u,v} \in \mathbb{R}^{d_e}$. A number of LGLMs have been developed to process graph-text associated data $\mathcal{D} = \{\mathcal{G}, c\}$, where $\boldsymbol{c} = [c_1, ..., c_{l_c}]$ is to the caption of the graph \mathcal{G} . For nodecentric tasks, \boldsymbol{c}_i will associate with the nodes (Tang et al., 2023), while in this paper we focus on graph-centric tasks, i.e., molecules and molecular captions (Liu et al., 2023c). Usually, an *l*-layer GNN is employed to encode a graph as:

$$\boldsymbol{h}_{u}^{(l)} = \text{COM}(\boldsymbol{h}_{u}^{(l-1)}, \text{AGG}(\{(\boldsymbol{h}_{u}^{(l-1)}, \boldsymbol{h}_{v}^{(l-1)}) | v \in \mathcal{N}(u)\}))$$
(1)

where $h_u^{(l)} \in \mathbb{R}^h$ refers to the node embedding of node u after l layers of GNN, AGG(·) is the aggregation function (e.g., mean) among the information from neighbors of node u, and COM is the operator for combining information of node u with its neighbors $\mathcal{N}(u)$ (e.g., concatenation). Then, after l message passing iterations, the graph-level embedding can be obtained as:

$$\boldsymbol{h}_{\mathcal{G}} = \operatorname{READOUT}\left(\{h_u^{(l)}|u \in \mathcal{V}\}\right),$$
 (2)

where READOUT(\cdot) is a pooling operator (e.g., mean pooling) among all the node embeddings. With the representations of the nodes and graphs, LGLMs can fuse the graph and language information in various ways, such as transforming into natural languages describing the graphs (Fatemi et al., 2024), or neural prompts within the LLMs (Tian et al., 2024). In addition, the embeddings can also be leveraged to post-process the LLM outputs (Liu et al., 2024b). Orthogonal to different fusion mechanisms, in this work, we focus on transforming graph embeddings into input tokens of LLMs, which can be formulated as (Tang et al., 2023; Chen et al., 2024; Liu et al., 2023; Cao et al., 2023; Li et al., 2024):

$$p_{\theta}(\boldsymbol{a}|\boldsymbol{q},\boldsymbol{h}) = \prod_{i=1}^{l_a} p_{\theta}(\boldsymbol{a}_i|\boldsymbol{q}, f_n(\boldsymbol{h}), \boldsymbol{a}_{< i}), \quad (3)$$

where the LGLM is required to approximate p_{θ} to output the desired answer a given the question q, and the graph tokens h adapted with adapter $f_n : \mathbb{R}^h \to \mathbb{R}^{h_e}$ that projects the graph tokens to the embedding space of LLMs. One could also incorporate the 1D information such as SMILES (Weininger, 1988) into q and a for alignment.

Molecular Foundation Models. There is a separate line of works aiming to develop language models for molecules and proteins - the language of lives, from 1D sequences such as SMILES (Irwin et al., 2022), 2D molecular graphs (Rong et al., 2020; Wang et al., 2022; Zhang et al., 2024a), 3D geometric conformations (Liu et al., 2022; Zhou et al., 2023), to scientific text (Beltagy et al., 2019) and multimodal molecule-text data (Liu et al., 2023b; Luo et al., 2023a; Christofidellis et al., 2023; Liu et al., 2024c; Su et al., 2022; Zeng et al., 2022; Srinivas & Runkana, 2024). The adopted backbones range from encoder-decoder architectures such as MolT5 (Edwards et al., 2022) and Galactica (Taylor et al., 2022), to auto-regressive language modeling (Luo et al., 2023b; Liu et al., 2023e). Inspired by the success of large vision-language models (Li et al., 2023a; Zhu et al., 2023; Liu et al., 2023a), the community further seeks to develop molecular foundation models built upon existing molecular

language models with more sophisticated graph information fusion modules. For example, Liu et al. (2023c); Zhao et al. (2023) develop advanced cross-modal adapters and generalized position embeddings to promote better alignment based on encoder-decoder-based molecular language models. Liang et al. (2023); Cao et al. (2023); Li et al. (2024) develop cross-modal adapters for decoder only language models such as Llama (Touvron et al., 2023a). Orthogonal to the aforementioned works, we focus more on *what information one shall extract from the molecules for the alignment*. We choose to build our methods upon decoderonly language models, with the hope of building a versatile agent that can perceive molecules beyond the language, image, and audio modalities (Xi et al., 2023).

In the meantime, existing works also try to enrich the molecule-language alignment with additional modalities, such as 2D (Liu et al., 2023c) and 3D (Li et al., 2024) information. In contrast, we focus on the intrinsic hierarchical information of the molecules, such as motifs.

Hierarchical Graph Representation Learning. The hierarchical nature has been widely incorporated in learning high-quality graph representations (Ying et al., 2018). Especially in molecular graphs, the high-order structural information naturally captures the existence of motifs and functional groups. Therefore, the hierarchy of atom-motif-molecule has been widely applied in self-supervised molecular representation learning (Zhang et al., 2021; Zang et al., 2023; Inae et al., 2023; Luong & Singh, 2023). Nevertheless, how to properly incorporate the hierarchical information in molecular instruction tuning of LGLMs remains unclear.

In addition, concurrent works by Park et al. (2024) and Hu & Li (2024) explored incorporating hierarchical graph information into LLMs. Nevertheless, they mostly focus on the architecture-level incorporation, while we show that it is crucial to integrate the hierarchical information in the instruction tuning data. More importantly, we highlight the consequences of inadequate alignment due to the lack of hierarchical information, i.e., hallucination, and demonstrate the usefulness of the hierarchical information in a wide range of downstream tasks.

3. Graph Tokenization in LGLMs

In this section, we analyze the limitations of node-centric tokenization, which is widely adopted in existing LGLMs.

3.1. Node-Centric Tokenization

Specifically, most existing LGLMs directly take the node tokens from GNNs as inputs to LLMs (Cao et al., 2023):

$$p_{\theta}(\boldsymbol{a}|\boldsymbol{q},\boldsymbol{h}) = \prod_{i=1}^{l_a} p_{\theta}(\boldsymbol{a}_i|\boldsymbol{q}, f_n(\boldsymbol{h}_1), ..., f_n(\boldsymbol{h}_n), \boldsymbol{a}_{< i}),$$
(4)

where $h_1, ..., h_n$ are node embeddings from a GNN typically pretrained through self-supervised learning on largescale molecular datasets such as ZINC250k (Sterling & Irwin, 2015), f_n is the adapter to project the node tokens to the LLM tokens. There are various options to tokenize a molecule (Liu et al., 2023d). In this work, we consider a state-of-the-art tokenizer (Xia et al., 2023) that pretrains a VQVAE (van den Oord et al., 2017) with masked atoms modeling and constructs a codebook \mathcal{Z} to discretize atoms: $z_u = \arg \min_i ||h_u - e_i||_2$, where $z_u \in \mathcal{Z}$ is the quantized index of atom u, and e_v is the codebook embedding of the *i*-th entry. The codebook is trained through a reconstruction loss with respect to some attribute v_i of atom *i*:

$$\mathcal{L}_{T} = \frac{1}{n} \sum_{i=1}^{n} (1 - \frac{\boldsymbol{v}_{i}^{T} \hat{\boldsymbol{v}}_{i}}{||\boldsymbol{v}_{i}|| \cdot || \hat{\boldsymbol{v}}_{i}||})^{\gamma} + \frac{1}{n} \sum_{i=1}^{n} ||\mathrm{sg}[\boldsymbol{h}_{i}] - \boldsymbol{e}_{z_{i}}||_{2}^{2} + \frac{\beta}{2} \sum_{i=1}^{n} ||\mathrm{sg}[\boldsymbol{e}_{z_{i}}] - \boldsymbol{h}_{i}||_{2}^{2},$$
(5)

where sg[·] is the stop-gradient operator in straight-through estimator (Bengio et al., 2013), \hat{v}_i is the reconstructed attribute of atom *i* with a decoder, and β is a hyperparamter. In Mole-BERT, the attribute is simply the type of atom. Mole-BERT also manually partitions the codebook into groups of common atoms such as carbon, nitrogen, and oxygen to avoid codebook conflicts (Xia et al., 2023).

Intuitively, the trained atom tokens encode some contextual information, such as the neighbors of the atoms. However, node-centric tokenization makes the molecule-language alignment more challenging, as LLMs have to additionally relate the multiple nodes to align the corresponding texts during the instruction tuning process. Specifically, in molecules, motifs or functional groups usually capture rich semantics, and often share many common atoms such as carbon, nitrogen, and oxygen (Bohacek et al., 1996). As shown in Fig. 2, both the carboxylic acid ("R-COOH") and the hydroperoxide ("R-OOH") functional groups all contain two oxygen atoms and a hydrogen atom. For a molecule with hydroperoxide attached to a scaffold with carbon atoms, it would be hard for LLMs to distinguish which functional group is present in the molecule. Furthermore, due to the loss of positional information in the node-centric tokenization (Liang et al., 2023; Cao et al., 2023), the limited expressivity of GNNs (Xu et al., 2019) and the positional biases of auto-regressive LLMs (Lu et al., 2022), it is more challenging for LLMs to relate the desired nodes in a motif, which will lead to subpar molecule-language alignment.

3.2. Motif Hallucination

To understand the issue of node-centric tokenization more clearly, we construct a simple benchmark called MotifHallu, to concretize the hallucination of common functional groups by LGLMs. Specifically, we consider the 38 common functional groups in RDKit¹ and leverage RDKit (Landrum, 2016) to detect the existence. We adopt 3,300 molecules from ChEBI-20 (Edwards et al., 2021) and query the existence of a functional group:

Then, we examine the outputs from LGLM meaning "Yes" or "No". For each molecule, we construct questions with positive answers for all kinds of functional groups detected in the molecule, and questions with negative answers for randomly sampled 6 functional groups from the remaining. Hence MotifHallu consists of 23, 924 questions. While it is easy to scale up MotifHallu with more molecules and functional groups, we find that the current scale is already sufficient to demonstrate the issue (Table 2).

4. Hierarchical Graph Tokenization

To improve the molecule-language alignment, we propose a new strategy called **HI**erarchical **GrapH Tokenization** (HIGHT), which contains a hierarchical graph tokenizer and a hierarchical molecular instruction tuning dataset to augment the inputs with hierarchical information.

4.1. Hierarchical Graph Tokenizer

Inspired by the success of hierarchical GNNs (Zhang et al., 2021; Zang et al., 2023), we transform the original molecular graph \mathcal{G} into a hierarchical graph \mathcal{G}' with motif-level and molecule-level nodes added in. Specifically, we leverage the Breaking of Retrosynthetically Interesting Chemical Substructures (BRICS) algorithm (Degen et al., 2008)² to detect and inject a set of k + 1 supernodes, denoted as $\mathcal{M} = {\mathcal{M}^{(1)}, ..., \mathcal{M}^{(k)}, \mathcal{M}^{(k+1)}}$, with k motifs and the original molecule $\mathcal{M}^{(k+1)} = \mathcal{G}$. Furthermore, denoting the set of nodes and edges in $\mathcal{M}^{(i)}$ as $\mathcal{V}_m^{(i)}$ and $\mathcal{E}_m^{(i)}$, respectively, we augment the original molecular graph \mathcal{G} as \mathcal{G}' with augmented nodes \mathcal{V}' and edges \mathcal{E}' :

$$\mathcal{V}' = \mathcal{V} \cup \{v_m^{(1)}, ..., v_m^{(k+1)}\}, \ \mathcal{E}' = \mathcal{E} \cup (\bigcup_{i=1}^{k+1} \mathcal{E}_{ma}^{(i)}), \ \ (6)$$

where $v_m^{(i)}$ is the motif super nodes added to the original molecule, and $\mathcal{E}_{ma}^{(i)} = \bigcup_{u \in \mathcal{V}_m^{(i)}} \{(u, v_m^{(i)})\}$ are the augmented edges connecting to the motif super node from nodes within the corresponding motif. We employ separate VQVAEs for atoms and motifs to learn meaningful code embeddings with

¹https://github.com/rdkit/rdkit/blob/ master/Data/FunctionalGroups.txt

²Note that HIGHT possesses a high degree of extensibility and can be augmented by incorporating advanced motif extraction techniques (such as (Zhang et al., 2021)).



Figure 2. Illustration of hallucination caused by node-centric tokenization. With only node-level tokens, LLMs have to relate the nodes within a specific functional group to align useful molecular structures with the corresponding language descriptions. Yet, due to the arbitrary order of atoms and position biases in LLMs, it is hard to recognize each functional group, leading to severe hallucinations.

several self-supervised learning tasks. The reconstructed attributes in Eq. 4 include atom types at the atom-level and the number of atoms at the motif-level (Zang et al., 2023).

Merely feeding the motif tokens with node tokens to LLMs still can not help distinguish the motifs from atoms properly, hence we propose to further attach positional encodings p to all of the tokens. We choose to use Laplacian positional embeddings (Dwivedi et al., 2020) while one could also adopt other variants (Ying et al., 2021). Since different types of tokens contain distinct semantics, we adopt separate adapters for different types of tokens. Denoting the motif tokens as $h_m^{(i)}$ for motif $\mathcal{M}^{(i)}$, generation with HIGHT is:

$$p_{\theta}(\boldsymbol{a}|\boldsymbol{q}, \boldsymbol{h}, \boldsymbol{h}_{m}) = \prod_{i=1}^{l_{a}} p_{\theta}(\boldsymbol{a}_{i}|\boldsymbol{q}, f_{n}(\boldsymbol{h}_{1}), ..., f_{n}(\boldsymbol{h}_{n}), \\ f_{m}(\boldsymbol{h}_{m}^{(1)}), ..., f_{g}(\boldsymbol{h}_{m}^{(k+1)}), \boldsymbol{a}_{< i}),$$
(7)

where $f_m(\cdot)$ and $f_g(\cdot)$ are the adapters for BRICS motifs and the original molecules, respectively.

4.2. Hierarchical Graph Instruction Tuning Dataset

Although HIGHT tokenizer properly extracts the hierarchical information from the input graph modality, it remains challenging to properly align the language to the corresponding molecular information, without the appearance of the respective captions in the texts. For example, if the caption does not contain any information about the water solubility of the hydroxide functional group ("-OH"), LGLMs will never know that "-OH" motif corresponds to the water solubility of the molecule, despite that HIGHT tokenizer extracts the "-OH" token. In fact, the commonly used molecular instruction tuning curated from PubChem (Kim et al., 2022) in existing LGLMs (Liu et al., 2023c; Cao et al., 2023; Li et al., 2024), contains surprisingly little information about motifs. Some samples are given in Appendix C.2. To this end, we propose HiPubChem, which augments the molecular instruction tuning dataset with captions of the functional groups. We consider both the positive and negative appearances of motifs: For the positive case, we directly append the caption of all functional groups detected with RDKit. We also include a brief introduction of the functional groups to provide fine-grained information for molecule-language alignment. For the negative case, we randomly sample k_{neg} motifs not appeared in the molecule to explicitly instruct LGLMs on the absence of the motifs. Despite the simple augmentation strategy, we find that HiPubChem significantly reduces the hallucination issue and improves the alignment performance.

4.3. Hierarchical Graph Instruction Tuning

We use a two-stage instruction tuning (Cao et al., 2023).

Stage 1 Alignment Pretraining. We curate a new moleculetext paired dataset from PubChem following the pipeline of Liu et al. (2023b). We set the cutoff date by Jan. 2024, and filter out unmatched pairs and low-quality data, which results in 295k molecule-text pairs. Furthermore, we construct the HiPubChem-295k dataset. The first stage mainly warms up the adapter to properly project the graph tokens with the LLM embedding space. To avoid feature distortion, both the LLM and the GNN encoder are frozen.

Stage 2 Task-specific Instruction Tunning. With a properly trained adapter, we further leverage the task-specific instruction tuning datasets from MoleculeNet (Wu et al., 2017), ChEBI-20 (Mendez et al., 2019), and Mol-Instructions (Fang et al., 2024). More details are given in Appendix C. In Stage 2, we still keep the GNN encoder frozen, while tuning both the adapter and the LLM (with low-rank adaptation, i.e., LoRA (Hu et al., 2022)).

Table 1. Detailed results in motif hallucinations on MotifHallu.	Due to the imbalance of samples from positive and negative classes
we incorporate diverse evaluation metrics to provide a detailed comp	parison between different methods in terms of hallucination.

Метнод	Macro F1 ↑	F1 (pos) ↑ 4, 124	F1 (neg) ↑ 19,800	Micro F1 ↑	AUROC ↑	Acc \uparrow	Yes Ratio
GIMLET (Zhao et al., 2023)	50.0	0.1	99.9	0.05	49.9	82.6	0.2
Galactica-6.7B (Taylor et al., 2022)	56.6	17.5	95.7	12.9	50.7	77.6	8.5
InstructMol (Cao et al., 2023)	52.6	95.7	9.5	28.3	48.4	20.0	94.5
HIGHT	66.8	85.5	48.2	29.7	53.2	39.1	69.4

Table 2. Results of motif hallucinations on MotifHallu.

Method	F1 (pos) ↑	F1 (neg) ↑	Acc \uparrow	Yes Ratio
Node-centric Tokenization				
InstructMol-G	95.7	9.5	19.9	94.5
InstructMol-G (LLama-2-7b-chat)	99.6	2.8	18.3	98.7
InstructMol-GS	97.1	10.6	20.9	94.4
Hierarchical Tokenization				
HIGHT-G	85.5	48.2	39.1	74.7
HIGHT-G (LLama-2-7b-chat)	55.1	65.2	46.6	49.3
HIGHT-GS	84.5	42.7	35.1	73.1
Ablation variants of HIGHT				
HIGHT-G w/o HiPubChem	96.6	12.5	21.6	96.6
HIGHT-GS w/o HiPubChem	98.2	6.5	19.4	93.3

5. Experimental Evaluation

We conduct extensive experiments to compare HIGHT with previous node-centric tokenization across 14 real-world tasks, including property prediction, molecular description, and chemical reaction prediction. The details and examples regarding the datasets and tasks involved in the experiments are given in Appendix C. We briefly introduce the setups below and leave the details in Appendix D.

5.1. Experimental settings

Architecture. The GNN backbone is a 5-layer GIN (Xu et al., 2019) with a hidden dimension of 300. The adapter is a single-layer MLP. We consider base LLMs of vicuna-v-1.3-7B (Chiang et al., 2023) for all the tasks and llama-2-7B-chat (Touvron et al., 2023b) for ablation studies.

Baselines. Since the focus of this work lies in the tokenization, our main comparison focuses on between HIGHT and node-centric tokenization. Nevertheless, we also include a series of existing LGLMs based on non-regression LLMs and regression LLMs, to provide an overview of the performance achieved by HIGHT. We would like to note that there are existing differences in pretraining data and information used between HIGHT and those baselines. For details, please refer to Table 7 in the Appendix.

For the node-centric based tokenization, we implement the baseline mainly based on InstructMol (Cao et al., 2023) with a VQVAE tokenizer from Mole-BERT (Xia et al., 2023). HIGHT is implemented based on the same architecture with only the tokenizer replaced. We use the suffix "-G" to refer to LGLMs with only 2D graph input and "-GS" to refer to LGLMs with both 2D graph and 1D selfies input (Krenn et al., 2019; Fang et al., 2024; Cao et al., 2023). We do

not include the baselines with "-GS" for tasks other than MotifHallu as we find that incorporating the 1D input does not always bring improvements in the experiments.

For non-regression-based models, including the pretrained models such as KV-PLM (Zeng et al., 2022), GraphCL (You et al., 2020) and GraphMVP (Liu et al., 2022), and molecular foundation models that are trained with tremendous molecule-centric datasets such as MoIT5-based methods (Edwards et al., 2022), Galactica (Taylor et al., 2022), MoMu (Su et al., 2022), MoIFM (Luo et al., 2023a), Uni-Mol (Zhou et al., 2023), MoIXPT (Liu et al., 2023e), GIT-Mol (Liu et al., 2024c), and BioMedGPT (Luo et al., 2023b). We adopt the results from the previous works Fang et al. (2024); Cao et al. (2023) if applicable.

For regression-based LGLMs, we consider LLMs such as ChatGPT (OpenAI, 2022), Llama (Touvron et al., 2023a) as well as instruction tuned LLMs such as Alpaca (Dubois et al., 2023), Baize (Xu et al., 2023), ChatGLM (Zeng et al., 2023) and Vicuna (Chiang et al., 2023). We also consider parameter-efficient finetuned LLMs using the backbone of llama2 (Touvron et al., 2023b) as done by Mol-Instructions (Fang et al., 2024).

5.2. Motif Hallucination

We begin with a proof-of-concept study with motif hallucination. We mainly compare LGLMs with node-centric to that with HIGHT tokenization with MotifHallu after stage 1 instruction tuning. For non-regression-based models, we include two state-of-the-art LGLMs GIMLET (Zhao et al., 2023) and Galactica (Taylor et al., 2022). We do not include the other regression-based models as we found they consistently answered "Yes", making a nuanced F1 comparison less informative for them. To avoid the issue of format following, we compare the loss values by feeding the answers of "Yes" and "No" to the corresponding LLM, calculating the language modeling losses, and taking the one from "Yes" and "No" with a lower loss as the answer.

Reduction of hallucination. Due to the class imbalance issue in MotifHallu, we first report comprehensive metrics in Table 1. It can be found that HIGHT maintains great balance for both positive and negative classes compared to baselines. Especially, in terms of macro F1 scores that are

Table 3. Results of molecular property prediction tasks (regression) on QM9. We report the result in MAE. \dagger : few-shot in-context learning (ICL) results from (Fang et al., 2024). $\Delta\epsilon$ refers to the HOMO-LUMO energy gap.

Method	$\rm HOMO\downarrow$	$LUMO\downarrow$	$\Delta \epsilon \downarrow$	Avg \downarrow
Alpaca [†] (Dubois et al., 2023)	-	-	-	322.109
Baize [†] (Xu et al., 2023)	-	-	-	261.343
LLama2-7B (Touvron et al., 2023b) (5-shot ICL)	0.7367	0.8641	0.5152	0.7510
Vicuna-13B (Chiang et al., 2023) (5-shot ICL)	0.7135	3.6807	1.5407	1.9783
Mol-Instruction (Fang et al., 2024)	0.0210	0.0210	0.0203	0.0210
InstructMol-G	0.0111	0.0133	0.0147	0.0130
HIGHT-G	0.0078	0.0086	0.0095	0.0086

averaged across classes, respectively, HIGHT demonstrates significant improvements up to 14%.

The results of the tokenization-focused comparison are given in Table 2. Following the practice in LVLMs, we present the F1 scores, accuracies, and the ratio that the model answers "Yes" (Li et al., 2023c). Given the imbalance of positive and negative samples, we separately report the F1 scores for different classes. It can be found that the LGLMs with node-centric tokenization consistently answer with "Yes" despite the absence of the corresponding functional groups. In contrast, HIGHT significantly reduces the worst class hallucination up to 40% in terms of F1 scores, and improves the accuracies up to 30%. The improvements are consistent and significant with both vicunna and llama2 LLM backbones.

Ablations with different inputs and LLM backbones. We also conduct simple ablation studies by additionally incorporating the 1D sequence inputs with SELFIES (Fang et al., 2024; Cao et al., 2023). Contrary to previous results that additionally feeding the 1D sequence always improves the performance of LGLMs, we find that the additional 1D sequence may increase the degree of the hallucination. We suspect that it could be caused by the extremely long sequences of the SELFIES (Krenn et al., 2019) that may distract the attention signals of LLMs. Nevertheless, HIGHT still suffers less from the distraction and performs better.

In addition, when without HiPubChem (or with the HIGHT architecture), LGLMs will still suffer the hallucination, due to the low quality of the instruction tuning data, demonstrating the necessity of both components of HIGHT.

5.3. Molecular-Centric Benchmarks

Molecular property prediction requires LGLMs to answer about particular properties given the molecule. We use 8 datasets BACE, BBBP, HIV, SIDER, ClinTox, MUV, and Tox21 from MoleculeNet, and CYP450 from GIM-LET (Zhao et al., 2023) to evaluate the classification performance with ROC-AUC. We also adopt the regression-based property prediction datasets from (Fang et al., 2024), where we evaluate several quantum chemistry measures such as HUMO, LUMO, and HUMO-LUMO gap (Ramakrishnan et al., 2014) via Mean Absolute Error (MAE).

The results of molecular property prediction are given in Table 3 and Table 4 for regression and classification, respectively. We can find that HIGHT always significantly boosts the performance in both types of tasks. Remarkably, in CYP450 (Zhao et al., 2023), HIGHT significantly outperforms the state-of-the-art model, demonstrating the advances of LGLM with hierarchical graph tokenization. Interestingly, Llama-2 (Touvron et al., 2023b) can match the state-of-the-art performance in HIV in a few-shot setting, while performing significantly worse in other datasets, for which we suspect some data contamination might exist.

Molecular description requires the LGLMs to generate a caption of the molecule. We adopt the widely used benchmark ChEBI-20 (Edwards et al., 2021) which evaluates the linguistic distances of the generated molecule captions of molecular characteristics such as structure, properties, biological activities etc.. We report the metrics of BLEU (Papineni et al., 2002), ROUGE (Lin, 2004) and Meteor (Banerjee & Lavie, 2005). The LGLMs are trained using the ChEBI-20 train split, selected according to the best training loss, and evaluated using the test split.

As shown in Table 5, HIGHT consistently brings significant improvements over LGLMs with node-centric tokenization. Nevertheless, compared to the molecular foundation models such as MoT5 (Edwards et al., 2022) pretrained on a significant amount of molecule-text related corpus, there remains a gap for regression-based LGLMs even with HIGHT. The gap calls for future investigations on how to incorporate HIGHT into the pretraining of the LGLMs properly.

Chemical reaction prediction requires the LGLMs to predict the results of the chemical reaction analysis, which are crucial for AI-aided drug discovery (Fang et al., 2024). Reagent prediction aims to predict the suitable reagents for a particular chemical reaction. Forward reaction prediction aims to predict the products of a chemical reaction, given the reactants and the reagents. Retrosynthesis prediction aims to predict the suitable reactants given a target product. The inputs and outputs for chemical reaction related tasks adopt the SELFIES (Krenn et al., 2019) as recommended by (Fang et al., 2024). We report both linguistic distance metrics such as BLEU (Papineni et al., 2002) and Levenshtein (Yujian & Bo, 2007), and molecular similarity measures such as similarity of the molecular fingerprints (Landrum, 2016).

As shown in Table 6, across all tasks in chemical reaction prediction, LGLMs with HIGHT consistently and significantly improve the performances compared to the nodecentric tokenization. Meanwhile, LGLMs with HIGHT achieve state-of-the-art results in several tasks and metrics, compared to other regression-based LGLMs that even incorporate a stronger LLM backbone such as Mol-Instruction,

Table 4. ROC-AUC Results of molecular property prediction tasks (classification) on MoleculeNet (Wu et al., 2017). Evaluation on InstructMol and HIGHT adopt the likelihood of the tokens of "Yes" and "No". Most of the instruction tuning datasets are from GIMLET (Zhao et al., 2023). SIDER and ClinTox are converted following the MoleculeNet task description.

			ē			*		
Method	BACE \uparrow	BBBP ↑	$HIV\uparrow$	SIDER \uparrow	ClinTox	MUV ↑	Tox21 ↑	CYP450↑
# MOLECULES	1,513	2,039	41,127	1,427	1,478	93,087	7,831	16,896
# TASKS	1	1	1	27	2	17	12	5
KV-PLM (Zeng et al., 2022)	78.5	70.5	71.8	59.8	84.3	61.7	49.2	59.2
GraphCL (You et al., 2020)	75.3	69.7	78.5	60.5	76.0	69.8	73.9	-
GraphMVP-C (Liu et al., 2022)	81.2	72.4	77.0	60.6	84.5	74.4	77.1	-
MoleculeSTM-G (Liu et al., 2023b)	80.8	70.0	76.9	61.0	92.5	73.4	76.9	-
MoMu (Su et al., 2022)	76.7	70.5	75.9	60.5	79.9	60.5	57.8	58.0
MolFM (Luo et al., 2023a)	83.9	72.9	78.8	64.2	79.7	76.0	77.2	-
Uni-Mol (Zhou et al., 2023)	85.7	72.9	80.8	65.9	91.9	82.1	78.1	-
Galactica-1.3B (Taylor et al., 2022)	57.6	60.4	72.4	54.0	58.9	57.2	60.6	46.9
Galactica-6.7B (Taylor et al., 2022)	58.4	53.5	72.2	55.9	78.4	-	63.9	-
Galactica-30B (Taylor et al., 2022)	72.7	59.6	75.9	61.3	82.2	-	68.5	-
Galactica-120B (Taylor et al., 2022)	61.7	66.1	74.5	63.2	82.6	-	68.9	-
GIMLET (Zhao et al., 2023)	69.6	59.4	66.2	-	-	64.4	61.2	71.3
LLama-2-7b-chat (4-shot) (Touvron et al., 2023b)	76.9	54.2	67.8	-	-	46.9	62.0	57.6
LLama-2-13b-chat (4-shot) (Touvron et al., 2023b)	74.7	52.8	72.4	-	-	47.9	57.5	55.6
InstructMol-G	64.3	48.7	50.2	51.0	50.0	50.0	59.0	59.1
HIGHT-G	77.1	61.8	63.3	58.8	55.3	51.1	67.4	80.5

Table 5. Results of molecular description generation task on the test split of ChEBI-20.

Model	BLEU-2↑	BLEU-4↑	ROUGE-1↑	ROUGE-2↑	ROUGE-L↑	METEOR↑
MoT5-base (Edwards et al., 2022)	0.540	0.457	0.634	0.485	0.568	0.569
MoMu (MolT5-base) (Su et al., 2022)	0.549	0.462	-	-	-	0.576
MolFM (MolT5-base) (Luo et al., 2023a)	0.585	0.498	0.653	0.508	0.594	0.607
MolXPT (Liu et al., 2023e)	0.594	0.505	0.660	0.511	0.597	0.626
GIT-Mol-graph (Liu et al., 2024c)	0.290	0.210	0.540	0.445	0.512	0.491
GIT-Mol-SMILES (Liu et al., 2024c)	0.264	0.176	0.477	0.374	0.451	0.430
GIT-Mol-(graph+SMILES) (Liu et al., 2024c)	0.352	0.263	0.575	0.485	0.560	0.430
Text+Chem T5-augm-base (Christofidellis et al., 2023)	0.625	0.542	0.682	0.543	0.622	0.648
GPT-3.5-turbo (10-shot MolReGPT) (Li et al., 2023b)	0.565	0.482	0.623	0.450	0.543	0.585
GPT-4-0314 (10-shot MolReGPT) (Li et al., 2023b)	0.607	0.525	0.634	0.476	0.562	0.610
GPT-3.5-turbo (zero-shot) (Li et al., 2023b)	0.103	0.050	0.261	0.088	0.204	0.161
BioMedGPT-10B (Luo et al., 2023b)	0.234	0.141	0.386	0.206	0.332	0.308
Mol-Instruction (Fang et al., 2024)	0.249	0.171	0.331	0.203	0.289	0.271
InstructMol-G	0.481	0.381	0.554	0.379	0.488	0.503
HIGHT-G	0.504	0.405	0.570	0.397	0.502	0.524

and additional information of SELFIES.

5.4. Empirical Analysis

Generalist capabilities. We follow the previous practice in training and evaluating generalist models (Liu et al., 2023a) and consider the two settings: a) As shown in Fig. 3(a), we first train the model with all chemical reaction prediction data by 3 epochs to elicit the format following and the knowledge adaption capabilities of the LGLMs after stage 1. The models are named with "-All"; b) As shown in Fig. 3(b), we train the model with retrosynthesis task data and evaluate the zero-shot transfer performance on forward reaction prediction. Under both settings, we can find that HIGHT boosts the generalist capabilities significantly.

Computation overhead. In Appendix E.1, we report the computation overhead of pretraining and inference as well as tunable parameters of HIGHT and InstructMol. It can be

found that, although HIGHT requires longer training time and relatively higher tunable parameters, the absolute values are not high. Moreover, during inference, as LLM latency consumes most of the computation, HIGHT can even reduce the inference latency by generating more concise answers.

Ablation studies. To better understand the effectiveness of distinct components in HIGHT, we conduct ablation studies that train InstructMol (Cao et al., 2023) with the laplacian positional encodings or with HiPubChem, as given in Fig. 3(c). We can find that, merely incorporating positional encoding or hierarchical instruction tuning is not sufficient to achieve the same performance as HIGHT. On the contrary, without a proper architecture design as HIGHT, LGLMs with previous node-centric tokenization with HiPubChem will confuse LLMs and even lead to degenerated downstream task performances. In addition, we also compare LGLMs with llama2 backbone. As shown in Fig. 3(a),

Table 6. Results of chemical reaction tasks. These tasks encompass reagent prediction, forward reaction prediction, and retrosynthes	is. †:
few-shot ICL results from (Fang et al., 2024). *: use task-specific instruction data to finetune.	

Model	Exact†	BLEU↑	Levenshtein ↓	RDK FTS↑	MACCS FTS↑	Morgan FTS↑	V ALIDITY [↑]
Reagent Prediction							
Alpaca [†] (Dubois et al., 2023)	0.000	0.026	29.037	0.029	0.016	0.001	0.186
Baize [†] (Xu et al., 2023)	0.000	0.051	30.628	0.022	0.018	0.004	0.099
ChatGLM [†] (Zeng et al., 2023)	0.000	0.019	29.169	0.017	0.006	0.002	0.074
Llama [†] (Touvron et al., 2023a)	0.000	0.003	28.040	0.037	0.001	0.001	0.001
Vicuna [†] (Chiang et al., 2023)	0.000	0.010	27.948	0.038	0.002	0.001	0.007
Mol-Instruction (Fang et al., 2024)	0.044	0.224	23.167	0.237	0.364	0.213	1.000
Llama-7b* (Touvron et al., 2023a)(LoRA)	0.000	0.283	53.510	0.136	0.294	0.106	1.000
InstructMol-G	0.031	0.429	31.447	0.389	0.249	0.220	1.000
HIGHT-G	0.050	0.462	28.970	0.441	0.314	0.275	1.000
Forward Reaction Prediction							
Alpaca [†] (Dubois et al., 2023)	0.000	0.065	41.989	0.004	0.024	0.008	0.138
Baize [†] (Xu et al., 2023)	0.000	0.044	41.500	0.004	0.025	0.009	0.097
ChatGLM [†] (Zeng et al., 2023)	0.000	0.183	40.008	0.050	0.100	0.044	0.108
Llama [†] (Touvron et al., 2023a)	0.000	0.020	42.002	0.001	0.002	0.001	0.039
Vicuna [†] (Chiang et al., 2023)	0.000	0.057	41.690	0.007	0.016	0.006	0.059
Mol-Instruction (Fang et al., 2024)	0.045	0.654	27.262	0.313	0.509	0.262	1.000
Llama-7b* (Touvron et al., 2023a)(LoRA)	0.012	0.804	29.947	0.499	0.649	0.407	1.000
InstructMol-G	0.031	0.853	24.790	0.512	0.362	0.303	0.993
HIGHT-G	0.037	0.869	23.759	0.590	0.394	0.340	0.993
Retrosynthesis							
Alpaca [†] (Dubois et al., 2023)	0.000	0.063	46.915	0.005	0.023	0.007	0.160
Baize [†] (Xu et al., 2023)	0.000	0.095	44.714	0.025	0.050	0.023	0.112
ChatGLM [†] (Zeng et al., 2023)	0.000	0.117	48.365	0.056	0.075	0.043	0.046
Llama [†] (Touvron et al., 2023a)	0.000	0.036	46.844	0.018	0.029	0.017	0.010
Vicuna [†] (Chiang et al., 2023)	0.000	0.057	46.877	0.025	0.030	0.021	0.017
Mol-Instruction (Fang et al., 2024)	0.009	0.705	31.227	0.283	0.487	0.230	1.000
Llama-7b* (Touvron et al., 2023a)(LoRA)	0.000	0.283	53.510	0.136	0.294	0.106	1.000
InstructMol-G	0.001	0.835	31.359	0.447	0.277	0.241	0.996
HIGHT-G	0.008	0.863	28.912	0 564	0 340	0.309	1.000



HIGHT still significantly boosts the performance. More ablation studies are provided in Appendix E.

6. Conclusions

This paper presents HIGHT, a novel hierarchical graph tokenization technique. By incorporating the hierarchical graph information, HIGHT improves the molecule-language alignment performance, reducing hallucinations and boosting accuracy in molecular tasks. Nevertheless, the current focus on molecular graphs requires further verification for wider applicability to other forms of graph data, such as those originating from social networks. Despite the limitation, HIGHT represents a significant step forward in advancing graph comprehension capability of LLMs, and highlighting paths for future research in this direction.

Meanwhile, incorporating 3D information into the graphlanguage alignment is also a promising future direction, especially for broader scientific tasks such as single-cell modeling and understanding. For example, built upon HIGHT, one could design a new 3D tokenizer to accommodate 3D properties of motifs, scale up 3D data to include amino acids in proteins and certain recurrent structures in RNA sequences, incorporate 3D positional encoding, and curate instruction tuning data with 3D descriptive captions.

Acknowledgments

We thank the reviewers for their valuable comments. JC was supported by RGC Young Collaborative Research Grant No. C2005-24Y.

Impact Statement

This paper mainly focuses on how to best represent graph information for LLMs to better understand the graphs. We demonstrate the effectiveness of our method on moleculecentric tasks, which could facilitate the broader use of LLMs for tasks like AI-aided drug discovery and human-machine interactions in biomedicine. Besides, this paper does not raise any ethical concerns. This study does not involve any human subjects, practices, to data set releases, potentially harmful insights, methodologies and applications, potential conflicts of interest and sponsorship, discrimination/bias/fairness concerns, privacy and security issues, legal compliance, and research integrity issues.

References

- Banerjee, S. and Lavie, A. 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*, pp. 65–72, 2005. (Cited on pages 7 and 24)
- Beltagy, I., Lo, K., and Cohan, A. SciBERT: A pretrained language model for scientific text. In *Conference on Empirical Methods in Natural Language Processing*, pp. 3615–3620, 2019. (Cited on pages 3 and 24)
- Bengio, Y., Léonard, N., and Courville, A. C. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint*, arXiv:1308.3432, 2013. (Cited on page 4)
- Bohacek, R. S., McMartin, C., and Guida, W. C. The art and practice of structure-based drug design: A molecular modeling perspective. *Medicinal Research Reviews*, 16 (1):3–50, 1996. (Cited on pages 1, 4 and 20)
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S. M., Nori, H., Palangi, H., Ribeiro, M. T., and Zhang, Y. Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv preprint, arXiv:2303.12712, 2023. (Cited on page 1)
- Cao, H., Liu, Z., Lu, X., Yao, Y., and Li, Y. Instructmol: Multi-modal integration for building a versatile and reliable molecular assistant in drug discovery, 2023. (Cited on pages 1, 3, 4, 5, 6, 7, 8 and 17)

- Chen, R., Zhao, T., Jaiswal, A., Shah, N., and Wang, Z. Llaga: Large language and graph assistant. *arXiv preprint*, arXiv:2402.08170, 2024a. (Cited on page 3)
- Chen, Y., Zhang, Y., Bian, Y., Yang, H., Ma, K., Xie, B., Liu, T., Han, B., and Cheng, J. Learning causally invariant representations for out-of-distribution generalization on graphs. In *Advances in Neural Information Processing Systems*, 2022. (Cited on page 17)
- Chen, Y., Bian, Y., Zhou, K., Xie, B., Han, B., and Cheng, J. Does invariant graph learning via environment augmentation learn invariance? In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. (Cited on page 17)
- Chen, Y., Bian, Y., Han, B., and Cheng, J. How interpretable are interpretable graph neural networks? In *International Conference on Machine Learning*, 2024b. (Cited on page 17)
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., Stoica, I., and Xing, E. P. Vicuna: An opensource chatbot impressing gpt-4 with 90%* chatgpt quality, 2023. URL https://lmsys.org/blog/ 2023-03-30-vicuna/. (Cited on pages 6, 7, 9, 22, 26 and 27)
- Christofidellis, D., Giannone, G., Born, J., Winther, O., Laino, T., and Manica, M. Unifying molecular and textual representations via multi-task language modelling. In *International Conference on Machine Learning*, volume 202, pp. 6140–6157, 2023. (Cited on pages 3, 8 and 17)
- Degen, J., Wegscheid-Gerlach, C., Zaliani, A., and Rarey, M. On the art of compiling and using 'drug-like' chemical fragment spaces. *ChemMedChem*, 3:1503–1507, 2008. (Cited on page 4)
- Doreian, P. and Woodard, K. L. Defining and locating cores and boundaries of social networks. *Social Networks*, 16 (4):267–293, 1994. (Cited on page 17)
- Dubois, Y., Li, X., Taori, R., Zhang, T., Gulrajani, I., Ba, J., Guestrin, C., Liang, P., and Hashimoto, T. B. Alpacafarm: A simulation framework for methods that learn from human feedback, 2023. (Cited on pages 6, 7, 9, 26 and 27)
- Durant, J. L., Leland, B. A., Henry, D. R., and Nourse, J. G. Reoptimization of mdl keys for use in drug discovery. *Journal of chemical information and computer sciences*, 42 6:1273–80, 2002. (Cited on page 24)
- Dwivedi, V. P., Joshi, C. K., Luu, A. T., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982*, 2020. (Cited on page 5)

- Edwards, C., Zhai, C., and Ji, H. Text2mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 595–607, 2021. (Cited on pages 4, 7, 21 and 24)
- Edwards, C., Lai, T. M., Ros, K., Honke, G., Cho, K., and Ji, H. Translation between molecules and natural language. In *Conference on Empirical Methods in Natural Language Processing*, pp. 375–413, 2022. (Cited on pages 3, 6, 7, 8 and 17)
- Fan, W., Wang, S., Huang, J., Chen, Z., Song, Y., Tang, W., Mao, H., Liu, H., Liu, X., Yin, D., and Li, Q. Graph machine learning in the era of large language models (llms). *arXiv preprint*, arXiv:2404.14928, 2024. (Cited on pages 1 and 2)
- Fang, Y., Liang, X., Zhang, N., Liu, K., Huang, R., Chen, Z., Fan, X., and Chen, H. Mol-instructions: A largescale biomolecular instruction dataset for large language models. In *International Conference on Learning Representations*, 2024. (Cited on pages 5, 6, 7, 8, 9, 20, 24, 26 and 27)
- Fatemi, B., Halcrow, J., and Perozzi, B. Talk like a graph: Encoding graphs for large language models. In *International Conference on Learning Representations*, 2024. (Cited on page 3)
- Hastings, J., Owen, G., Dekker, A., Ennis, M., Kale, N., Muthukrishnan, V., Turner, S., Swainston, N., Mendes, P., and Steinbeck, C. Chebi in 2016: Improved services and an expanding collection of metabolites. *Nucleic Acids Research*, 44:D1214 – D1219, 2015. (Cited on pages 18 and 21)
- Hu, C. and Li, H. Exploring hierarchical molecular graph representation in multimodal llms. *arXiv preprint*, arXiv:2411.04708, 2024. (Cited on page 3)
- Hu, E. J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. (Cited on page 5)
- Inae, E., Liu, G., and Jiang, M. Motif-aware attribute masking for molecular graph pre-training. In *NeurIPS* 2023 Workshop: New Frontiers in Graph Learning, 2023. (Cited on pages 2 and 3)
- Irwin, R., Dimitriadis, S., He, J., and Bjerrum, E. J. Chemformer: a pre-trained transformer for computational chemistry. *Machine Learning Science Technology*, 3(1): 15022, 2022. (Cited on page 3)

- Jin, B., Liu, G., Han, C., Jiang, M., Ji, H., and Han, J. Large language models on graphs: A comprehensive survey. *arXiv preprint*, arXiv:2312.02783, 2023. (Cited on pages 1 and 2)
- Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B. A., Thiessen, P. A., Yu, B., Zaslavsky, L., Zhang, J., and Bolton, E. E. PubChem 2023 update. *Nucleic Acids Research*, 51(D1):D1373– D1380, 10 2022. (Cited on pages 5 and 17)
- Krenn, M., Hase, F., Nigam, A., Friederich, P., and Aspuru-Guzik, A. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1, 2019. (Cited on pages 6, 7 and 20)
- Landrum, G. Rdkit: Open-source cheminformatics software, 2016. URL https://github.com/rdkit/ rdkit/releases/tag/Release_2016_09_4. (Cited on pages 4, 7, 21 and 24)
- Li, J., Li, D., Savarese, S., and Hoi, S. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, pp. 19730–19742, 2023a. (Cited on pages 1 and 3)
- Li, J., Liu, Y., Fan, W., Wei, X.-Y., Liu, H., Tang, J., and Li, Q. Empowering molecule discovery for moleculecaption translation with large language models: A chatgpt perspective. *arXiv preprint arXiv:2306.06615*, 2023b. (Cited on page 8)
- Li, S., Liu, Z., Luo, Y., Wang, X., He, X., Kawaguchi, K., Chua, T.-S., and Tian, Q. Towards 3d moleculetext interpretation in language models. In *International Conference on Learning Representations*, 2024. (Cited on pages 3 and 5)
- Li, Y., Du, Y., Zhou, K., Wang, J., Zhao, W. X., and Wen, J. Evaluating object hallucination in large vision-language models. *arXiv preprint*, arXiv:2305.10355, 2023c. (Cited on pages 7, 21, 22 and 24)
- Li, Y., Li, Z., Wang, P., Li, J., Sun, X., Cheng, H., and Yu, J. X. A survey of graph meets large language model: Progress and future directions. *arXiv preprint*, arXiv:2311.12399, 2023d. (Cited on pages 1 and 2)
- Liang, Y., Zhang, R., Zhang, I., and Xie, P. Drugchat: Towards enabling chatgpt-like capabilities on drug molecule graphs. *arXiv preprint*, arXiv:2309.03907, 2023. (Cited on pages 3 and 4)
- Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, 2004. (Cited on pages 7 and 24)

- Liu, C., Chen, Y., Liu, T., Gong, M., Cheng, J., Han, B., and Zhang, K. Discovery of the hidden world with large language models. In *Advances in Neural Information Processing Systems*, pp. 102307–102365, 2024a. (Cited on page 17)
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. In *NeurIPS*, 2023a. (Cited on pages 1, 3 and 8)
- Liu, H., Feng, J., Kong, L., Liang, N., Tao, D., Chen, Y., and Zhang, M. One for all: Towards training one graph model for all classification tasks. In *International Conference* on Learning Representations, 2024b. (Cited on page 3)
- Liu, P., Ren, Y., Tao, J., and Ren, Z. Git-mol: A multi-modal large language model for molecular science with graph, image, and text. *Computers in Biology and Medicine*, pp. 108073, 2024c. (Cited on pages 3, 6 and 8)
- Liu, S., Wang, H., Liu, W., Lasenby, J., Guo, H., and Tang, J. Pre-training molecular graph representation with 3d geometry. In *International Conference on Learning Representations*, 2022. (Cited on pages 3, 6, 8 and 17)
- Liu, S., Nie, W., Wang, C., Lu, J., Qiao, Z., Liu, L., Tang, J., Xiao, C., and Anandkumar, A. Multi-modal molecule structure-text model for text-based retrieval and editing. *Nature Machine Intelligence*, 5(12):1447–1457, 2023b. (Cited on pages 3, 5, 8 and 17)
- Liu, Z., Li, S., Luo, Y., Fei, H., Cao, Y., Kawaguchi, K., Wang, X., and Chua, T.-S. MolCA: Molecular graphlanguage modeling with cross-modal projector and unimodal adapter. In *Conference on Empirical Methods in Natural Language Processing*, 2023c. (Cited on pages 3, 5 and 17)
- Liu, Z., Shi, Y., Zhang, A., Zhang, E., Kawaguchi, K., Wang, X., and Chua, T.-S. Rethinking tokenizer and decoder in masked graph modeling for molecules. In *Advances in Neural Information Processing Systems*, 2023d. (Cited on page 4)
- Liu, Z., Zhang, W., Xia, Y., Wu, L., Xie, S., Qin, T., Zhang, M., and Liu, T.-Y. MolXPT: Wrapping molecules with text for generative pre-training. In *Annual Meeting of the Association for Computational Linguistics*, pp. 1606– 1616. Association for Computational Linguistics, 2023e. (Cited on pages 3, 6, 8 and 17)
- Lu, Y., Bartolo, M., Moore, A., Riedel, S., and Stenetorp, P. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pp. 8086–8098, 2022. (Cited on page 4)

- Luo, Y., Yang, K., Hong, M., Liu, X. Y., and Nie, Z. Molfm: A multimodal molecular foundation model, 2023a. (Cited on pages 3, 6, 8 and 17)
- Luo, Y., Zhang, J., Fan, S., Yang, K., Wu, Y., Qiao, M., and Nie, Z. Biomedgpt: Open multimodal generative pre-trained transformer for biomedicine, 2023b. (Cited on pages 3, 6 and 8)
- Luong, K.-D. and Singh, A. Fragment-based pretraining and finetuning on molecular graphs. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. (Cited on pages 2 and 3)
- Mao, H., Chen, Z., Tang, W., Zhao, J., Ma, Y., Zhao, T., Shah, N., Galkin, M., and Tang, J. Graph foundation models. *arXiv preprint*, arXiv:2402.02216, 2024. (Cited on pages 1 and 2)
- Mendez, D., Gaulton, A., Bento, A. P., Chambers, J., Veij, M. D., Felix, E., Magariños, M. P., Mosquera, J. F., Mutowo-Meullenet, P., Nowotka, M., Gordillo-Marañón, M., Hunter, F. M. I., Junco, L., Mugumbate, G., Rodríguez-López, M., Atkinson, F., Bosc, N., Radoux, C. J., Segura-Cabrera, A., Hersey, A., and Leach, A. R. Chembl: towards direct deposition of bioassay data. *Nucleic Acids Research*, 47(Database-Issue):D930–D940, 2019. (Cited on page 5)
- Miao, S., Liu, M., and Li, P. Interpretable and generalizable graph learning via stochastic attention mechanism. *International Conference on Machine Learning*, 2022. (Cited on page 17)
- Milo, R., Shen-Orr, S. S., Itzkovitz, S., Kashtan, N., Chklovskii, D. B., and Alon, U. Network motifs: simple building blocks of complex networks. *Science*, 298 5594: 824–7, 2002. (Cited on page 1)
- OpenAI. Chatgpt. https://chat.openai.com/ chat/, 2022. (Cited on pages 1 and 6)
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics*, 2002. (Cited on pages 7 and 24)
- Park, J., Bae, M., Ko, D., and Kim, H. J. LLamo: Large language model-based molecular graph assistant. In *Annual Conference on Neural Information Processing Systems*, 2024. (Cited on page 3)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in*

Neural Information Processing Systems, pp. 8024–8035, 2019. (Cited on page 24)

- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners, 2019. (Cited on page 1)
- Ramakrishnan, R., Dral, P. O., Dral, P. O., Rupp, M., and von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 2014. (Cited on page 7)
- Ribeiro, P., Paredes, P., Silva, M. E. P., Aparicio, D., and Silva, F. A survey on subgraph counting: Concepts, algorithms, and applications to network motifs and graphlets. *ACM Computing Survey*, 54(2), 2021. (Cited on page 17)
- Rong, Y., Bian, Y., Xu, T., Xie, W., WEI, Y., Huang, W., and Huang, J. Self-supervised graph transformer on largescale molecular data. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 12559–12571. Curran Associates, Inc., 2020. (Cited on page 3)
- Schneider, N., Sayle, R. A., and Landrum, G. A. Get your atoms in order - an open-source implementation of a novel and robust molecular canonicalization algorithm. *Journal of chemical information and modeling*, 55 10: 2111–20, 2015. (Cited on page 24)
- Srinivas, S. S. and Runkana, V. Crossing new frontiers: Knowledge-augmented large language model prompting for zero-shot text-based de novo molecule design. *arXiv preprint*, arXiv:2408.11866, 2024. (Cited on page 3)
- Sterling, T. and Irwin, J. J. Zinc 15 ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337, 2015. (Cited on pages 1 and 4)
- Su, B., Du, D., Yang, Z.-Q., Zhou, Y., Li, J., Rao, A., Sun, H., Lu, Z., and rong Wen, J. A molecular multimodal foundation model associating molecule graphs with natural language. *arXiv preprint arXiv:2209.05481*, 2022. (Cited on pages 3, 6, 8 and 17)
- Tang, J., Yang, Y., Wei, W., Shi, L., Su, L., Cheng, S., Yin, D., and Huang, C. Graphgpt: Graph instruction tuning for large language models. *arXiv preprint*, arXiv:2310.13023, 2023. (Cited on pages 2 and 3)
- Taylor, R., Kardas, M., Cucurull, G., Scialom, T., Hartshorn,
 A. S., Saravia, E., Poulton, A., Kerkez, V., and Stojnic,
 R. Galactica: A large language model for science. *arXiv* preprint, arXiv:2211.09085, 2022. (Cited on pages 3, 6, 8 and 17)

- Tian, Y., Song, H., Wang, Z., Wang, H., Hu, Z., Wang, F., Chawla, N. V., and Xu, P. Graph neural prompting with large language models. In *Thirty-Eighth AAAI Conference on Artificial Intelligence*, pp. 19080–19088, 2024. (Cited on page 3)
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models. *arXiv preprint*, arXiv:2302.13971, 2023a. (Cited on pages 1, 3, 6, 9 and 26)
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Canton-Ferrer, C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. Llama 2: Open foundation and finetuned chat models. arXiv preprint, arXiv:2307.09288, 2023b. (Cited on pages 6, 7, 8 and 27)
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017. (Cited on page 4)
- Wang, Q., Lin, Y., Chen, Y., Schmidt, L., Han, B., and Zhang, T. Do CLIP models always generalize better than imagenet models? In *Annual Conference on Neural Information Processing Systems*, 2024. (Cited on page 17)
- Wang, Y., Wang, J., Cao, Z., and Farimani, A. B. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3):279–287, 2022. (Cited on page 3)
- Wei, L., Gao, J., Zhao, H., and Yao, Q. Towards versatile graph learning approach: from the perspective of large language models. *arXiv preprint*, arXiv:2402.11641, 2024. (Cited on pages 1 and 2)
- Weininger, D. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28:31–36, 1988. (Cited on page 3)

- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. S. Moleculenet: A benchmark for molecular machine learning. arXiv preprint arXiv:1703.00564, 2017. (Cited on pages 5, 8, 20 and 24)
- Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Cheng, W., Zhang, Q., Qin, W., Zheng, Y., Qiu, X., Huang, X., and Gui, T. The rise and potential of large language model based agents: A survey, 2023. (Cited on page 3)
- Xia, J., Zhao, C., Hu, B., Gao, Z., Tan, C., Liu, Y., Li, S., and Li, S. Z. Mole-BERT: Rethinking pre-training graph neural networks for molecules. In *International Conference on Learning Representations*, 2023. (Cited on pages 4, 6, 17 and 22)
- Xu, C., Guo, D., Duan, N., and McAuley, J. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. *arXiv preprint*, arXiv:2304.01196, 2023. (Cited on pages 6, 7, 9, 26 and 27)
- Xu, J., Chen, Y., Dong, X., Lan, M., Huang, T., Bian, Q., Cheng, J., and Ke, Y. Brainood: Out-of-distribution generalizable brain network analysis. *arXiv preprint*, arXiv:2502.01688, 2025. (Cited on page 17)
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference* on *Learning Representations*, 2019. (Cited on pages 1, 4, 6 and 22)
- Yao, T., Chen, Y., Chen, Z., Hu, K., Shen, Z., and Zhang, K. Empowering graph invariance learning with deep spurious infomax. In *Forty-first International Conference* on Machine Learning, 2024. (Cited on page 17)
- Yao, T., Chen, Y., Hu, K., Liu, T., Zhang, K., and Shen, Z. Learning graph invariance by harnessing spuriosity. In *International Conference on Learning Representations*, 2025. (Cited on page 17)
- Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., and Liu, T.-Y. Do transformers really perform badly for graph representation? In *Advances in Neural Information Processing Systems*, 2021. (Cited on page 5)
- Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., and Leskovec, J. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems*, 2018. (Cited on pages 3 and 17)

- You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., and Shen, Y. Graph contrastive learning with augmentations. In *Advances in Neural Information Processing Systems*, pp. 5812–5823, 2020. (Cited on pages 6, 8 and 17)
- Yu, J., Xu, T., Rong, Y., Bian, Y., Huang, J., and He, R. Graph information bottleneck for subgraph recognition. In *International Conference on Learning Representations*, 2021. (Cited on page 17)
- Yujian, L. and Bo, L. A normalized levenshtein distance metric. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):1091–1095, 2007. (Cited on pages 7 and 24)
- Zang, X., Zhao, X., and Tang, B. Hierarchical molecular graph self-supervised learning for property prediction. *Communications Chemistry*, 6(1):34, 2023. (Cited on pages 2, 3, 4, 5 and 22)
- Zeng, A., Liu, X., Du, Z., Wang, Z., Lai, H., Ding, M., Yang, Z., Xu, Y., Zheng, W., Xia, X., Tam, W. L., Ma, Z., Xue, Y., Zhai, J., Chen, W., Liu, Z., Zhang, P., Dong, Y., and Tang, J. GLM-130b: An open bilingual pre-trained model. In *International Conference on Learning Representations*, 2023. (Cited on pages 6, 9 and 26)
- Zeng, Z., Yao, Y., Liu, Z., and Sun, M. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications*, 13(862), 2022. (Cited on pages 3, 6, 8 and 17)
- Zhang, J., Bian, Y., Chen, Y., and Yao, Q. Unimot: Unified molecule-text language model with discrete token representation. *arXiv preprint arXiv:2408.00863*, 2024a. (Cited on page 3)
- Zhang, J., Huang, J., Jin, S., and Lu, S. Vision-language models for vision tasks: A survey. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 2024b. (Cited on page 1)
- Zhang, Z., Liu, Q., Wang, H., Lu, C., and Lee, C.-K. Motifbased graph self-supervised learning for molecular property prediction. In *Advances in Neural Information Processing Systems*, pp. 15870–15882, 2021. (Cited on pages 2, 3 and 4)
- Zhang, Z., Liu, Q., Hu, Q., and Lee, C.-K. Hierarchical graph transformer with adaptive node sampling. *arXiv* preprint arXiv:2210.03930, 2022. (Cited on page 17)
- Zhao, H., Liu, S., Ma, C., Xu, H., Fu, J., Deng, Z.-H., Kong, L., and Liu, Q. GIMLET: A unified graph-text model for instruction-based molecule zero-shot learning. In *Neural Information Processing Systems*, 2023. (Cited on pages 1, 3, 6, 7, 8, 17 and 19)

- Zhou, G., Gao, Z., Ding, Q., Zheng, H., Xu, H., Wei, Z., Zhang, L., and Ke, G. Uni-mol: A universal 3d molecular representation learning framework. In *The Eleventh International Conference on Learning Representations*, 2023. (Cited on pages 3, 6, 8 and 17)
- Zhu, D., Chen, J., Shen, X., Li, X., and Elhoseiny, M. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. (Cited on page 3)

Appendix of HIGHT

Contents

A	More Future Works	17						
B	Comparison between other LGLMs 17							
С	Details of Instruction Tuning Datasets	17						
	C.1 Details of the PubChem Dataset	17						
	C.2 Details of HiPubChem Dataset	18						
	C.3 Details of Property Prediction Dataset	20						
	C.4 Details of Reaction Prediction Dataset	20						
	C.5 Details of Molecular Description Dataset	21						
	C.6 Details of MotifHallu Dataset	21						
D	Details of Experiments	22						
E	More Ablation Studies	24						
	E.1 Computation overhead	24						
	E.2 Ablation studies with different setups of the tokenizers	24						

A. More Future Works

Built upon HIGHT, there are several promising future directions. For example, one could extend this study to more types of graphs, such as social networks and knowledge graphs, by exploring the crucial substructures therein:

- Indeed, motifs generically exist in other types of graphs and are crucial for a variety of tasks (Ribeiro et al., 2021). For example, cliques can define boundaries between groups of people in social networks (Doreian & Woodard, 1994). The idea of HIGHT could be seamlessly applied to other graphs where we have some prior knowledge about critical motifs.
- Meanwhile, when we do not have prior knowledge about the motifs, the GNNs intrinsically model the hierarchical nature of graphs in different orders (Ying et al., 2018) and thus can be integrated into LGLMs to learn the hierarchical graph information. A similar idea has been verified successful in graph transformers (Zhang et al., 2022).
- Furthermore, one could also adopt interpretable GNNs to identify the critical subgraphs for the task (Yu et al., 2021; Miao et al., 2022; Chen et al., 2024b) that capture the underlying causal information about the underlying tasks (Chen et al., 2022; 2023; Yao et al., 2024; Liu et al., 2024a; Yao et al., 2025; Xu et al., 2025). It is also interesting to further investigate the hallucinations caused by the spurious correlations during the alignment (Wang et al., 2024).

B. Comparison between other LGLMs

Table 7. Comparison between other LGLMs in terms of the backbone, instruction tuning, downstream usage for Molecular Property prediction, and capable tasks. It can be found that HIGHT is capable of various tasks, given limited pre-training data and information. Note that compared to the instruction tuning data for other LGLMs, such as KV-PLM (Zeng et al., 2022), which consists of papers with detailed information about molecules, the text descriptions in HIGHT contain relatively simple sentences.

Model	Backbone	Information	Instruction Tuning Data	Downstream	# Tasks
HIGHT/InstructMol (Cao et al., 2023)	GNN+Llama	2D Graph+Text	HiPubChem-295k	LoRA	7
GraphCL (You et al., 2020)	GNN	2D Graph	Downstream training data	Finetuning	1
GraphMVP (Liu et al., 2022)	GNN	2D+3D graph	Geom-50K	Finetuning	1
MoleculeSTM (Xia et al., 2023)	GNN+SciBERT	2D Graph+3D Graph+Text	PubChemSTM-280K	Finetuning	6
KV-PLM (Zeng et al., 2022)	BERT	1D SMILES+Text	S2orc-300K academic papers	Finetuning	5
MolT5 (Edwards et al., 2022)	T5	1D SMILES+Text	C4+ZINC(100M)	Finetuning	2
Text+Chem T5 (Christofidellis et al., 2023)	T5	1D SMILES+Text	Multi task-33.5M	Finetuning	5
MoMu (Su et al., 2022)	GNN+BERT	2D Graph+Text	Graph-Docuemnt Pair-15.6K	Finetuning	4
MolFM (Luo et al., 2023a)	GNN+BERT	Knowledge Graph+2D+3D Graph+Text	KG-15K+S2ORC-37M	Finetuning	4
Uni-Mol (Zhou et al., 2023)	Transformer	3D Graph	molecule-209M+protein-3.2M	Finetuning	4
Galactica (Taylor et al., 2022)	GPT	1D SMILES/Text	Documents-59M+chemicla prompts-2.5M	zero-shot	12
MolXPT (Liu et al., 2023e)	GPT2	1D SMILES+Text	Mixed text-68M	Finetuning	2
GIMLET (Zhao et al., 2023)	GNN+T5	2D Graph+Text	ChemBL-730K	zero-shot	2

C. Details of Instruction Tuning Datasets

We provide a summary of the datasets for instruction tuning and evaluation in this paper as in Table 8. Meanwhile, we also list the data sources and the corresponding licenses of the sources for each task and dataset. Then, we will elaborate more on the details of the datasets in the following subsections.

C.1. Details of the PubChem Dataset

PubChem³ is one of the largest public molecule database (Kim et al., 2022), and has been widely adopted by the alignment training of LGLMs (Liu et al., 2023c;b; Cao et al., 2023). Our construction of PubChem predominantly follows Liu et al. (2023b). We will briefly describe the main steps and interested readers may refer the details to (Liu et al., 2023b):

- We curate the data from PubChem using the official API and set the data cutoff date as 12 Jan. 2024. It downloads both the molecular structure (e.g., SMILES, 2D molecular graphs) in SDF format, and the text descriptions.
- Then, we will filter out molecules that do not have descriptions or can not match via the PubChem ID. In the descriptions, the molecule names are replaced with "This molecule", in order to facilitate LLMs to understand the instructions.

³https://pubchem.ncbi.nlm.nih.gov

Datasets	Train	Test	Content
PubChem	295,228	N/A	Molecules and the associated descriptions from PubChem.
HiPubChem	295,228	N/A	Molecules and the associated descriptions from PubChem and about functional groups in the molecule.
MoleculeNet-HIV	32,901	4,113	Question answering about the ability of the molecule to inhibit HIV replication.
MoleculeNet-BACE	1,210	152	Question answering about the ability of the molecule to bind to the BACE1 protein
MoleculeNet-BBBP	1,631	204	Question answering about the ability of the molecule to diffuse across the brain blood barrier.
MoleculeNet-SIDER	1,141	143	Question answering about the ability of the side effects.
MoleculeNet-ClinTox	1,188	148	Question answering about the toxicology.
MoleculeNet-MUV	74,469	9,309	Question answering about PubChem bioAssay
MoleculeNet-Tox21	6,877	860	Question answering about Toxicology in the 21st century
CYP45-	13,516	1,690	Question answering about CYP PubChem BioAssay CYP 1A2, 2C9, 2C19, 2D6, 3A4 inhibition.
Property Prediction (Regression)	360,113	1,987	Question answering about the quantum mechanics proper- ties of the molecule.
Forward Reaction Prediction	124,384	1,000	Question answering about the products of a chemical re- action, given specific reactants and reagents.
Reagent Prediction	124,384	1,000	Question answering about suitable catalysts, solvents, or ancillary substances required for a specific chemical reac- tion.
Retrosynthesis Prediction	128,684	1,000	Question answering about the reactants and reagents of a chemical reaction, given specific products.
ChEBI-20	26,407	3,300	Molecules and the associated Chemical Entities of Biolog- ical Interest (ChEBI) (Hastings et al., 2015) annotations.
MotifHallu	N/A	23,924	Question answering about existing functional groups in the molecule.

Table 8. Summary of datasets involved in our paper.

Finally, the curation generates 295k molecule-text pairs that we term as PubChem-295k. PubChem-295k will be mainly used for the stage 1 alignment training.

C.2. Details of HiPubChem Dataset

HiPubChem augments the molecular instruction tuning dataset with captions of the functional groups. We consider both the positive and negative appearances of motifs when augmenting the instructions. For the positive case, we directly append the caption of all functional groups detected with RDKit:

This molecule has <#> of <functional group name> groups.

For the negative case, we randomly sample k_{neg} that do not appear in the molecule:

This molecule has no <functional group name> groups.

Tasks/Datasets	Data Sources	License URL	License Note
PubChem, HiPubChem	PubChem	https://www.nlm.nih. gov/web_policies.html	Works produced by the U.S. gov- ernment are not subject to copyright protection in the United States. Any such works found on National Li- brary of Medicine (NLM) Web sites may be freely used or reproduced without permission in the U.S.
Reaction Prediction	USPTO	https://www.uspto.gov/ learning-and-resources/ open-data-and-mobility	It can be freely used, reused, and redistributed by anyone.
Property Prediction	MoleculeNet	https://opensource.org/ license/mit/	Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associ- ated documentation files (the "Soft- ware"), to deal in the Software with out restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so
Property Prediction	CYP450	https://www.nlm.nih. gov/web_policies.html	The data is from Zhao et al. (2023) that curates PubChem BioAssay CYP 1A2, 2C9, 2C19, 2D6, 3A4 inhibition. Thus it shares the same license as PubChem.
Molecular Description, MotifHallu	ChEBI	https://creativecommons. org/licenses/by/4.0/	You are free to: Share — copy and redistribute the material in any medium or format. Adapt — remix, transform, and build upon the mate- rial for any purpose, even commer- cially.

Table 9. Summary of data resources and licenses of datasets involved in our paper.

Table 10. Summary of inputs and outputs of the tasks in experiments.

	input	output
motif hallucination	molecule and question about the existence of a motif	yes or no
molecular property prediction (classification)	molecule and question about the existence of the property	yes or no
molecular property prediction (regression)	molecule and question about the value of the property	property value
molecular caption	molecule and question asking for the molecular caption	molecular caption
chemical reaction prediction	molecules and question about the reaction	molecular results

Table 11. Examples of PubChem and HiPubChem datasets.

PubChem	HiPubChem
SMILES: $CC(=O)OC(CC(=O)[O-])C[N+](C)(C)C$	
This molecule is an O-acylcarnitine having acetyl as the acyl substituent.	This molecule has 1 carboxylic acids functional group. This molecule
It has a role as a human metabolite. It is functionally related to an acetic	has no methyl amide, or amide, or nitro or thiols groups. This molecule
acid. It is a conjugate base of an O-acetylcarnitinium.	is an O-acylcarnitine having acetyl as the acyl substituent. It has a role
	as a human metabolite. It is functionally related to an acetic acid. It is a
	conjugate base of an O-acetylcarnitinium.
SMILES: CCN(CC)CCOC(=0)C(Cc1cccc2cccc12)CC1CCCO1	
This molecule is a member of naphthalenes.	This molecule has 0 functional groups. This molecule is a member of naphthalenes.
SMILES: Cc1c2[nH]c(c1CCC(=O)O)Cc1[nH]c(c(CCC(=O)O)c1C)Cc1[[nH]c(c(CCC(=O)O)c1C)Cc1[nH]c(c(C)c1CCC(=O)O)C2
This molecule is a coproporphyrinogen. It has a role as an Escherichia	This molecule has 1 carboxylic acids functional groups. This molecule
coli metabolite and a mouse metabolite. It is a conjugate acid of a	has no methyl amide, or diazo, or cyano or thiols groups. This molecule
coproporphyrinogen III(4-).	is a coproporphyrinogen. It has a role as an Escherichia coli metabolite
	and a mouse metabolite. It is a conjugate acid of a coproporphyrinogen
	III(4-).

Despite the simple augmentation strategy, we find that HiPubChem significantly reduces the hallucination issue, and improves the molecule-language alignment performance.

For comparison, we provide examples of PubChem and HiPubChem in Table 11.

.

C.3. Details of Property Prediction Dataset

The task of molecular property prediction mainly aims to predict certain biochemical or physical properties of molecules. Usually, these properties have a close relation with the molecular substructures (i.e., functional groups) (Bohacek et al., 1996). In this work, we consider the scenarios of both binary classification based and the regression based molecular property prediction, and the datasets are mainly derived from MoleculeNet (Wu et al., 2017).

For the classification, we consider three subtasks, HIV, BACE, and BBBP. The HIV subtask mainly evaluates whether the molecule is able to impede the replication of the HIV virus. The BACE subtask mainly evaluates the binding capability of a molecule to the BACE1 protein. The BBBP subtask mainly evaluates the capability of a molecule to passively diffuse across the human brain blood barrier. For task-specific instruction tuning, we convert those classification based datasets into instructions. Examples are given in Table 12.

Dataset	Question	Answer
HIV	SMILES: N=C1OC2(c3ccccc3)C3=C(OC(=NC)N2C)C(=O)OC3(c2cccc2)N1C	
	Please help me evaluate whether the given molecule can impede the replication of the HIV virus.	No
BACE	SMILES: CN(C(=O)CCc1cc2cccc2nc1N)C1CCCCC1	
	Can the given molecule bind to the BACE1 protein?	Yes
BBBP	$SMILES: \ Cc1c[nH+][o+]c(C([NH])CC(C)C(C)(C)N(C(C)(C)C)C(C)(N)N)c1[O-]$	
	Can the given molecule passively diffuse across the brain blood barrier?	Yes

Table 12. Examples of the property prediction (classification) datasets.

Table 13. Examples of the property prediction (regression) datasets.

Question	Answer
SELFIES: [0][=C][0][C][C][C][C][Ring1][=Branch1][C][Ring1][Ring2]	
Can you give me the energy difference between the HOMO and LUMO orbitals of this molecule?	0.2756
SELFIES: [C][C][=Branch1][C][=O][N][Branch1][C][C][C][=Branch1][C][=O][N]	
What is the lowest unoccupied molecular orbital (LUMO) energy of this molecule?	-0.0064
SELFIES: [C][C][=C][O][C][=C][Ring1][Branch1][C][Branch1][C][C][C]	
Please provide the highest occupied molecular orbital (HOMO) energy of this molecule.	-0.2132

For regression, we adopt the instruction tuning data from Mol-Instructions (Fang et al., 2024). The regression based property prediction focuses on predicting the quantum mechanics properties of the molecules. The 1D sequence information in this task is given by SELFIES (Krenn et al., 2019). The original data is sourced from the QM9 subset of the MolculeNet (Wu et al., 2017). There are three subtasks: (i) Highest occupied molecular orbital (HOMO) energy prediction; (ii) Lowest occupied molecular orbital (LUMO) energy prediction; (iii) and HUMO-LUMO gap energy prediction. Some examples of the regression based property prediction dataset are given in Table 13.

C.4. Details of Reaction Prediction Dataset

We adopt three chemical reaction related tasks from Mol-Instructions (Fang et al., 2024): Forward reaction prediction, reagent prediction, and retrosynthesis prediction. The input and output contain 1D sequence information given by SELFIES (Krenn et al., 2019). Some examples of the Mol-Instructions datasets are given in Table 14, where the SELFIES represented molecules are denoted as "¡SELFIES¿" for clarity.

The task of forward reaction prediction aims to predict the possible products of a chemical reaction. The input includes the SELFIES sequences of the reactant and reagent of the chemical reaction. And the model needs to predict the SELFIES of the products. The original data is sourced from USPTO ⁴, which consists of chemical reactions of organic molecules extracted from American patents and patent applications.

The task of reagent reaction prediction aims to predict the suitable catalysts, solvents, and ancillary substances with respect to a chemical reaction. The input includes the SELFIES sequences of the chemical reaction. The original data is sourced

⁴https://developer.uspto.gov/data

Task	Examples
Forward Reaction Prediction	Question: With the provided reactants and reagents, propose a potential product. SELFIES
	Answer: ¡SELFIES¿
Reagent Prediction	<i>Question:</i> Please suggest some possible reagents that could have been used in the following chemical reaction. The reaction is <i>¡SELFIES;</i>
	Answer: ¡SELFIES¿
Retrosynthesis Prediction	Question: Please suggest potential reactants for the given product. The product is: ¡SELFIES;
	Answer: ;SELFIES;

from USPTO ⁵, as the other tasks.

The task of retrosynthesis prediction aims to reverse engineer a particular compound by predicting the potential reactants or reagents that are required to synthesis the compound. The input includes the SELFIES sequences of the target product. The original data is sourced from USPTO 6 , similar to the other tasks.

C.5. Details of Molecular Description Dataset

For the molecular description task, we adopt a widely used dataset ChEBI-20 (Edwards et al., 2021). Based on the molecules from PubChem, Edwards et al. (2021) collected the Chemical Entities of Biological Interest (ChEBI) (Hastings et al., 2015) annotations of the molecules, which are the descriptions of molecules. We transform the task into the instructions, and present some samples in Table 15. The authors collect 33, 010 molecule-text pairs and split them into training (80%), validation (10%), and testing (10%) subsets. We mainly adopt the original training split to tune the model and evaluate the tuned model on the original test split.

Table 15. Examples of the molecular description datasets.			
Question	Answer		
SMILES: C1=CC=C(C=C1)[As](=O)(O)[O-]			
Could you give me a brief overview of this molecule?	The molecule is the organoarsonic acid anion formed by loss of a single proton from the arsonic acid grouping in phenylarsonic acid. It is a conjugate base of a phenylar- sonic acid.		
<i>SMILES: CCCCCCCCCC(=0)OC(=0)CCCCCCCCCC</i>			
Could you provide a description of this molecule?	The molecule is an acyclic carboxylic anhydride resulting from the formal condensation of the carboxy groups of two molecules of dodecanoic acid. It derives from a dode- canoic acid.		
SMILES: CCCCNC=O			
Please give me some details about this molecule.	The molecule is a member of the class of formamides that is formamide substituted by a butyl group at the N atom. It has a role as a human metabolite. It derives from a formamide.		

Table 15. Examples of the molecular description datasets.

C.6. Details of MotifHallu Dataset

The MotifHallu is mainly used to measure the hallucination of common functional groups by LGLMs. For the construction of MotifHallu, we consider the common functional groups in RDKit⁷ as shown in Table 16. There are 39 common functional groups, while we neglect the one with the name of "???".

Then, we leverage RDKit (Landrum, 2016) to detect the existence of the left 38 valid functional groups within a molecule. We consider 3, 300 molecules from ChEBI-20 test split (Edwards et al., 2021), and adopt the query style as for large vision-language models (Li et al., 2023c) that queries the existence of specific functional group one by one:

⁵https://developer.uspto.gov/data

⁶https://developer.uspto.gov/data

⁷https://github.com/rdkit/rdkit/blob/master/Data/FunctionalGroups.txt

Is there a <functional group name> in the molecule?

Examples of MotifHallu are given in Table 17.

During the evaluation, we detect whether the LGLM gives outputs meaning "Yes" or "No" following the practice in (Li et al., 2023c). For each molecule, we construct questions with positive answers for all kinds of functional groups detected in the molecule, and questions with negative answers for randomly sampled 6 functional groups from the 38 common functional groups in RDKit. The construction finally yields 23,924 query answer pairs about the existence of functional groups in the molecule. While it is easy to scale up MotifHallu by automatically considering more molecules and a broader scope of functional groups, we find that the current scale is already sufficient to demonstrate the hallucination phenomena in LGLMs.

D. Details of Experiments

Implementation of graph tokenizer. We implement the GNN tokenizer/encoder based on the same GNN backbone, which is a 5-layer GIN (Xu et al., 2019). The hidden dimension is 300. For the node-centric tokenization, we employ the VQVAE GNN tokenizer from Mole-BERT (Xia et al., 2023) and adopt self-supervised learning tasks from the official Mole-BERT implementation.⁸ For HIGHT, we train the VQVAE with the self-supervised learning tasks from (Zang et al., 2023) based on the official implementation.⁹ Meanwhile, we set the hyperparameters of GNN tokenizer training the same as those recommended by (Xia et al., 2023; Zang et al., 2023).

After training the tokenizer, we adopt the GNN encoder within the tokenizer instead of the codebook embeddings as we empirically find that the GNN embeddings perform better than that using the VQVAE codebook embeddings.

Implementation of LGLMs. For the cross-modal adapters, we implement it as a single-layer MLP with an input dimension of 300 as our main focus is the tokenization. For HIGHT, we adopt three distinct adapters to handle the node-level, motif-level and graph-level embeddings. Meanwhile, we also adopt a Laplacian position encodings with respect to the supernode-augmented graphs. The dimension of the Laplacian position encoding is set to 8, therefore the input dimensions of the adapters in HIGHT will be 308.

For the LoRA adapters, we use a LoRA rank of 128 and a scaling value α of 256 for molecular property prediction (classification) in order to better fit with the task, and use a LoRA rank of 64 and a scaling value α of 16 for all the remaining methods and tasks.

For the base LLM, we mainly adopt vicuna-v-1.3-7B (Chiang et al., 2023). The overall scale of parameters is around 6.9B.

Implementation of instruction tuning. In stage 1 instruction tuning, we train all methods based on PubChem-295k dataset. The training goes 5 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. If there is an OOM issue, we will decrease the batch size a little bit to 40. The learning rate is set to 2×10^{-3} for all methods.

For classification-based property prediction, the training goes 20 epochs, with a batch size of 128 (distributed to 4 GPUs) by default. If there is an OOM issue, we will decrease the batch size a little bit to 64. The learning rate is set to 8×10^{-5} for all methods.

For regression-based property prediction, the training goes 5 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. The learning rate is set to 2×10^{-5} for all methods.

For molecular description, the training goes 50 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. If there is an OOM issue, we will decrease the batch size a little bit to 32. The learning rate is set to 8×10^{-5} for all methods.

For forward reaction prediction, the training goes 5 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. The learning rate is set to 2×10^{-5} for all methods.

For reagent prediction, the training goes 5 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. The learning rate is set to 2×10^{-5} for all methods.

⁸https://github.com/junxia97/Mole-BERT

⁹https://github.com/ZangXuan/HiMol

Chemical Representation	SMARTS	Name	
-NC(=O)CH3	*-[N;D2]-[C;D3](=O)-[C;D1;H3]	methyl amide	
-C(=O)O	*-C(=O)[O;D1]	carboxylic acids	
-C(=O)OMe	*-C(=O)[O;D2]-[C;D1;H3]	carbonyl methyl ester	
-C(=O)H	*-C(=O)-[C;D1]	terminal aldehyde	
-C(=O)N	*-C(=O)-[N;D1]	amide	
-C(=O)CH3	*-C(=O)-[C;D1;H3]	carbonyl methyl	
-N=C=O	*-[N;D2]=[C;D2]=[O;D1]	isocyanate	
-N=C=S	*-[N;D2]=[C;D2]=[S;D1]	isothiocyanate	
	Nitrogen containing groups		
-NO2	*-[N;D3](=[O;D1])[O;D1]	nitro	
-N=O	*-[N;R0]=[O;D1]	nitroso	
=N-O	*=[N;R0]-[O;D1]	oximes	
=NCH3	*=[N;R0]-[C;D1;H3]	Imines	
-N=CH2	*-[N;R0]=[C;D1;H2]	Imines	
-N=NCH3	*-[N;D2]=[N;D2]-[C;D1;H3]	terminal azo	
-N=N	*-[N;D2]=[N;D1]	hydrazines	
-N#N	*-[N;D2]#[N;D1]	diazo	
-C#N	*-[C;D2]#[N;D1]	cyano	
	S containing groups		
-SO2NH2	*-[S;D4](=[O;D1])(=[O;D1])-[N;D1]	primary sulfonamide	
-NHSO2CH3	*-[N;D2]-[S;D4](=[O;D1])(=[O;D1])-[C;D1;H3]	methyl sulfonamide	
-SO3H	*-[S;D4](=O)(=O)-[O;D1]	sulfonic acid	
-SO3CH3	*-[S;D4](=O)(=O)-[O;D2]-[C;D1;H3]	methyl ester sulfonyl	
-SO2CH3	*-[S;D4](=O)(=O)-[C;D1;H3]	methyl sulfonyl	
-SO2Cl	*-[S;D4](=O)(=O)-[C1]	sulfonyl chloride	
-SOCH3	*-[S;D3](=O)-[C;D1]	methyl sulfinyl	
-SCH3	*-[S;D2]-[C;D1;H3]	methylthio	
-S	*-[S;D1]	thiols	
=S	*=[S;D1]	thiocarbonyls	
	Miscellaneous fragments		
-X	*-[#9,#17,#35,#53]	halogens	
-tBu	*-[C;D4]([C;D1])([C;D1])-[C;D1]	t-butyl	
-CF3	*-[C;D4](F)(F)F	trifluoromethyl	
-C#CH	*-[C;D2]#[C;D1;H]	acetylenes	
-cPropyl	*-[C;D3]1-[C;D2]-[C;D2]1	cyclopropyl	
	Teeny groups		
-OEt	*-[O;D2]-[C;D2]-[C;D1;H3]	ethoxy	
-OMe	*-[O;D2]-[C;D1;H3]	methoxy	
-0	*-[O;D1]	side-chain hydroxyls	
=O	*=[O;D1]	side-chain aldehydes or ketones	
-N	*-[N;D1]	primary amines	
=N	*=[N;D1]	???	
#N	*#[N:D1]	nitriles	

Table 16. List of functional groups from RDKit used to construct MotifHallu. The functional group with the name "???" is neglected.

For retrosynthesis prediction, the training goes 5 epochs, with a batch size of 64 (distributed to 4 GPUs) by default. The learning rate is set to 2×10^{-5} for all methods.

Training and evaluation. Throughout the paper, we use a max token length of 2048. Meanwhile, we adopt an AdamW optimizer with a warmup ratio of 3% for optimizing all models. We select the final model according to the best training loss.

Table 17. Examples of the MotifHallu dataset.

Question	Answer
SMILES: COC1=CC=CC2=C1C(=CN2)C/C(=N/OS(=O)(=O)[O-])/S[C@H]3[C@@H]([C@H]([C@H]([C@B])([CB])([H](O3)CO)O)O)O
Is there a methyl ester sulfonyl group in the molecule?	No
<i>SMILES: CN</i> (<i>C</i>) <i>C</i> (= <i>O</i>) <i>C</i> (<i>CCN1CCC</i> (<i>CC1</i>)(<i>C</i> 2= <i>CC</i> = <i>C</i> (<i>C</i> = <i>C</i> 2) <i>C</i> 1) <i>O</i>)(<i>C</i> 3= <i>CC</i> = <i>CC</i> = <i>C</i> 3) <i>C</i> 4= <i>CC</i> = <i>CC</i> = <i>C</i> 4	
Is there a carbonyl methyl ester group in the molecule?	Yes
<i>SMILES: CN</i> (<i>C</i>) <i>C</i> (= <i>O</i>) <i>C</i> (<i>CCN1CCC</i> (<i>CC1</i>)(<i>C</i> 2= <i>CC</i> = <i>C</i> (<i>C</i> = <i>C</i> 2) <i>C</i> 1) <i>O</i>)(<i>C</i> 3= <i>CC</i> = <i>CC</i> = <i>C</i> 3) <i>C</i> 4= <i>CC</i> = <i>CC</i> = <i>C</i> 4	
Is there a terminal aldehyde group in the molecule?	No

For the evaluation of classification-based property prediction, we adopt the ROC-AUC following the common practice (Wu et al., 2017).

For the evaluation of regression-based property prediction, we adopt the Mean Absolute Error (MAE) following the common practice (Fang et al., 2024).

For the evaluation of molecular description, we adopt BLEU-2, BLEU-4, ROUGE-1, ROUGE-2, ROUGE-L, and METEOR following the common practice (Papineni et al., 2002; Lin, 2004; Edwards et al., 2021). To improve the reliability of the evaluation, the metrics are computed based on the tokenizer scibert_scivocab_uncased of SciBERT (Beltagy et al., 2019).

We follow the common practice to evaluate models for the tasks of chemical reaction predictions (Fang et al., 2024). We adopt linguistic metrics such as BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004), METEOR (Banerjee & Lavie, 2005) and Levenshtein scores (Yujian & Bo, 2007). Meanwhile, we also validate the validity of the generated molecular sequences with RDKit (Landrum, 2016). In addition, several molecular similarity measures are also leveraged. Specifically, we present the MAE of the RDKit, MACCS, and Morgan fingerprints to assess the semantic similarity of the generated compounds and the ground truth ones (Durant et al., 2002; Schneider et al., 2015).

As for the MotifHallu, in order to avoid the drawbacks that LGLMs may output answers that do not follow the instructions, we compare the loss values by feeding the answers of "Yes" and "No", and take the one with a lower autoregressive language modeling loss as the answer. Following the practice in LVLMs, we present the F1 scores, accuracies, and the ratio that the model answers "Yes" (Li et al., 2023c). Given the severe imbalance of positive and negative samples, we separately report the F1 scores for positive and negative classes.

Software and hardware. We implement our methods with PyTorch 11.3 (Paszke et al., 2019). We run experiments on Linux Servers with NVIDIA V100 and NVIDIA A100 (40G) graphics cards with CUDA 11.7.

E. More Ablation Studies

E.1. Computation overhead

Table 18. Training Computational Overhead. We count the average graph size of PubChem and HiPubChem, where HiPubChem adds 9 additional tokens on average. The real preprocessing time and training time are shown below, which are estimated based on 4 A100 40G GPUs. Although HIGHT requires more time to train, the absolute computational overhead of HIGHT is not high.

	Graph Size	Preprocessing Time	Training Time
PubChem HiPubChem	34.39 43.21	16min 32sec	8hour 17min 59sec
nirubChem	43.21	2511111 55Sec	15Hour Sommi 25sec

E.2. Ablation studies with different setups of the tokenizers

In Table 21, we present more results of the ablation studies with different setups of HIGHT and node-centric tokenizer.

Table 19. Inference Computational Overhead. The inference computational overhead is estimated based on 4 A100 40G GPUs. During the inference, the LLM latency takes up the majority of time. A well-trained LGLM with HIGHT is able to generate more concise and valid answers and thus may take less time during inference.

	Property Prediction	MolCaption	Reagent Prediction	Forward Reaction	Retrosynthesis
InstructMol	14min 54sec	6hour 22min 27sec	56min 56sec	1hour 34min 28sec	1hour 50min 47sec
HIGHT	15min 12sec	4hour 59min 50sec	50min 29sec	1hour 22min 08sec	1hour 49min 42sec

Table 20. Number of Tunable Parameters during Training. When pretraining the GNN tokenizer, the number of tunable parameters is the number of parameters in GNN encoder; In stage 1, the number of tunable parameters is the number of parameters in the projector; In stage 2, the number of tunable parameters is the number of parameters in the projector; In stage 2, the number of tunable parameters is the number of parameters in LoRA.

	graph token dimension	GNN encoder	params in projector	params in tokenizer	LoRA
InstructMol	300d	1,860,905	1,232,896	3,093,801	159,907,840
HIGHT	300d	1,865,105	3,796,992	5,662,097	159,907,840

Table 21. More results of chemical reaction tasks with ablation studies. These tasks encompass reagent prediction, forward reaction prediction, and retrosynthesis. †: few-shot ICL results from (Fang et al., 2024). *: use task-specific instruction data to finetune.

Model	Exact↑	BLEU↑	Levenshtein↓	RDK FTS↑	MACCS FTS↑	Morgan FTS \uparrow	VALIDITY↑
Reagent Prediction							
Alpaca [†] (Dubois et al., 2023)	0.000	0.026	29.037	0.029	0.016	0.001	0.186
Baize [†] (Xu et al., 2023)	0.000	0.051	30.628	0.022	0.018	0.004	0.099
ChatGLM [†] (Zeng et al. 2023)	0.000	0.019	29 169	0.017	0.006	0.002	0.074
LLama ^{\dagger} (Touvron et al. 2023a)	0.000	0.003	28 040	0.037	0.001	0.001	0.001
Vicuna [†] (Chiang et al. 2023)	0.000	0.010	27 948	0.038	0.002	0.001	0.007
Mol-Instruction (Fanglet al. 2024)	0.000	0.224	23.167	0.030	0.364	0.213	1,000
LLama-7b [*] (Touvron et al., 2023a)(LoRA)	0.000	0.283	53.510	0.136	0.294	0.106	1.000
InstructMol-G	0.031	0.429	31,447	0.389	0.249	0.220	1.000
+Positional Encoding	0.009	0.423	30.833	0.370	0.231	0.197	0.986
+HiPubChem	0.016	0.473	30.455	0.369	0.237	0.194	0.990
+Large Tokenizer	0.040	0.454	29.163	0.416	0.284	0.248	1.000
InstructMol-GS	0.057	0.439	29.757	0.437	0.314	0.271	0.999
InstructMol+LLama-2-7b-chat	0.016	0.459	29.238	0.359	0.225	0.189	0.988
HIGHT-G	0.050	0.462	28,970	0.441	0.314	0.275	1.000
HIGHT-GS	0.067	0.482	27.167	0.462	0.346	0.303	1.000
HIGHT +LLama-2-7b-chat	0.057	0.495	26.591	0.453	0.333	0.293	1.000
Forward Reaction Prediction							
Alpaca [†] (Dubois et al., 2023)	0.000	0.065	41,989	0.004	0.024	0.008	0.138
$Baize^{\dagger}$ (Xu et al., 2023)	0.000	0.044	41.500	0.004	0.025	0.009	0.097
ChatGLM ^{\dagger} (Zeng et al. 2023)	0.000	0.183	40.008	0.050	0.100	0.044	0.108
L = 2023a	0.000	0.020	42 002	0.001	0.002	0.001	0.039
Vicuna [†] (Chiang et al. 2023)	0.000	0.020	41.600	0.007	0.016	0.001	0.059
Mol-Instruction (Fang et al. 2024)	0.000	0.654	27 262	0.313	0.010	0.000	1.000
LL ama-7b* (Touvron et al. 2023a)(LoRA)	0.043	0.804	29.947	0.499	0.649	0.202	1.000
InstructMol-G	0.012	0.853	24 790	0.512	0.362	0.303	0.993
+Positional Encoding	0.0102	0.829	26.622	0.419	0.328	0.268	0.981
+HiPubChem	0.011	0.819	26.010	0.396	0.315	0.264	0.975
+Large Tokenizer	0.040	0.861	24.051	0.544	0.380	0.328	0.996
InstructMol-GS	0.252	0.926	17.773	0.755	0.599	0.543	1.000
InstructMol+LLama-2-7b-chat	0.020	0.841	25.109	0.426	0.339	0.284	0.998
HIGHT-G	0.037	0.869	23.759	0.590	0.394	0.340	0.993
HIGHT-GS	0.293	0.935	16.687	0.774	0.618	0.566	1.000
HIGHT +LLama-2-7b-chat	0.042	0.873	23.854	0.590	0.402	0.344	0.996
Retrosynthesis							
Alpaca [†] (Dubois et al., 2023)	0.000	0.063	46.915	0.005	0.023	0.007	0.160
Baize [†] (Xu et al., 2023)	0.000	0.095	44.714	0.025	0.050	0.023	0.112
ChatGLM [†] (Zeng et al., 2023)	0.000	0.117	48.365	0.056	0.075	0.043	0.046
LLama [†] (Touvron et al., 2023a)	0.000	0.036	46.844	0.018	0.029	0.017	0.010
Vicuna [†] (Chiang et al. 2023)	0.000	0.057	46 877	0.025	0.030	0.021	0.017
Mol-Instruction (Fang et al., 2024)	0.009	0.705	31.227	0.283	0.487	0.230	1.000
LLama-7b* (Touvron et al., 2023a)(LoRA)	0.000	0.283	53.510	0.136	0.294	0.106	1.000
InstructMol-G	0.001	0.835	31.359	0.447	0.277	0.241	0.996
+Positional Encoding	0.000	0.793	33.859	0.295	0.218	0.192	0.983
+HiPubChem	0.000	0.755	35.811	0.282	0.218	0.177	0.997
+Large Tokenizer	0.001	0.842	30.613	0.459	0.287	0.263	0.999
InstructMol-GS	0.172	0.911	20.300	0.765	0.615	0.568	1.000
InstructMol+LLama-2-7b-chat	0.000	0.806	32.128	0.292	0.234	0.202	0.985
HIGHT-G	0.008	0.863	28.912	0.564	0.340	0.309	1.000
HIGHT-GS	0.202	0.914	20.194	0.772	0.623	0.577	0.999
HIGHT +LLama-2-7b-chat	0.006	0.865	28.964	0.563	0.338	0.306	0.999

Table 22. Full results of motif hallucinations on MotifHallu with ablation studies.

Method	F1 (pos) \uparrow	F1 (neg) ↑	F1 (avg) †
Node-centric Tokenization			
InstructMol-G	95.7	9.5	52.6
InstructMol-G+LLama-2-7b-chat	99.6	2.8	51.2
InstructMol-GS	97.1	10.6	53.8
Hierarchical Tokenization			
HIGHT-G	85.5	48.2	66.9
HIGHT-G+LLama-2-7b-chat	55.1	65.2	60.2
HIGHT-GS	84.5	42.7	63.6
Ablation variants			
InstructMol-G + Positional Encoding	96.4	19.8	58.1
<pre>InstructMol-G + HiPubChem</pre>	96.6	12.5	54.6
HIGHT-G w/o HiPubChem	96.6	12.5	54.6
HIGHT-GS w/o HiPubChem	98.2	6.5	52.4

Table 23. Results of molecular property prediction tasks (regression) on QM9 with ablation studies. We report the result in MAE. \dagger : few-shot in-context learning (ICL) results from (Fang et al., 2024). $\Delta \epsilon$ refers to the HOMO-LUMO energy gap.

Method	$\mathrm{HOMO}\downarrow$	LUMO \downarrow	$\Delta \epsilon \downarrow$	Avg ↓
LLM Based Generalist Models				
Alpaca [†] (Dubois et al., 2023)	-	-	-	322.109
Baize [†] (Xu et al., 2023)	-	-	-	261.343
LLama2-7B (Touvron et al., 2023b) (5-shot ICL)	0.7367	0.8641	0.5152	0.7510
Vicuna-13B (Chiang et al., 2023) (5-shot ICL)	0.7135	3.6807	1.5407	1.9783
Mol-Instruction (Fang et al., 2024)	0.0210	0.0210	0.0203	0.0210
InstructMol-G	0.0111	0.0133	0.0147	0.0130
+Positional Encodings	0.0300	0.0395	0.0357	0.0350
+HiPubChem	0.0305	4.4019	0.0494	1.1226
HIGHT-G	0.0078	0.0086	0.0095	0.0086