

# MolEdit: Knowledge Editing for Multimodal Molecule Language Models

Zhenyu Lei  
University of Virginia  
Charlottesville, USA  
vjd5zr@virginia.edu

Yinhan He  
University of Virginia  
Charlottesville, USA  
nee7ne@virginia.edu

Patrick Soga  
University of Virginia  
Charlottesville, USA  
zqe3cg@virginia.edu

Yushun Dong  
Florida State University  
Tallahassee, USA  
yushun.dong@fsu.edu

Yaochen Zhu  
University of Virginia  
Charlottesville, USA  
uqp4qh@virginia.edu

Jundong Li  
University of Virginia  
Charlottesville, USA  
jundong@virginia.edu

## Abstract

Understanding and continuously refining multimodal molecular knowledge is crucial for advancing biomedicine, chemistry, and materials science. Molecule language models (MoLMs) have become powerful tools in these domains, integrating structural representations (e.g., SMILES strings, molecular graphs) with rich contextual descriptions (e.g., physicochemical properties, biomedical applications). However, MoLMs can encode and propagate inaccuracies due to outdated web-mined training corpora or malicious manipulation, jeopardizing downstream discovery pipelines. While knowledge editing has been explored for general-domain AI, its application to MoLMs remains uncharted, presenting unique challenges due to the multifaceted and interdependent nature of molecular knowledge. In this paper, we take the first step toward MoLM editing for two critical tasks: molecule-to-caption generation and caption-to-molecule generation. To address molecule-specific challenges, we propose MolEdit, a powerful framework that enables targeted modifications while preserving unrelated molecular knowledge. MolEdit combines a Multi-Expert Knowledge Adapter that routes edits to specialized experts for different molecular facets with an Expertise-Aware Editing Switcher that activates the adapters only when input closely matches the stored edits across all expertise, minimizing interference with unrelated knowledge. To systematically evaluate editing performance, we introduce MEBench, a comprehensive benchmark assessing multiple dimensions, including Reliability (accuracy of the editing), Locality (preservation of irrelevant knowledge), and Generality (robustness to reformed queries). Across extensive experiments on two popular MoLM backbones, MolEdit delivers up to 18.8% higher Reliability and 12.0% better Locality than state-of-the-art editing baselines while maintaining efficiency. Our findings chart a clear path toward safer, continuously updatable scientific foundation models. The code is available at: <https://github.com/LzyFischer/MolEdit>.

## CCS Concepts

• Computing methodologies → Machine learning.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WSDM '26, Boise, ID, USA*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2292-9/2026/02  
<https://doi.org/10.1145/3773966.3777928>

## Keywords

Molecule Language Model, Multimodal, Knowledge Editing

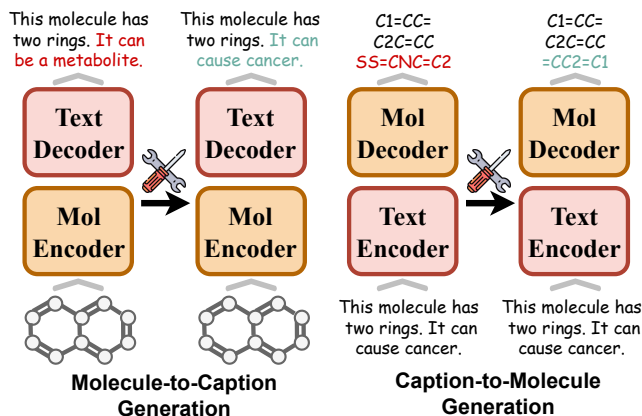
### ACM Reference Format:

Zhenyu Lei, Patrick Soga, Yaochen Zhu, Yinhan He, Yushun Dong, and Jundong Li. 2026. MolEdit: Knowledge Editing for Multimodal Molecule Language Models. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3773966.3777928>

## 1 Introduction

Understanding and continuously refining molecular knowledge is crucial across various scientific fields, such as biomedicine [45, 63], chemistry [26, 56], and materials science [21]. Pre-trained and fine-tuned on diverse multimodal data sourced from web such as PubChem [13], molecule language models (MoLMs) encode multimodal molecular knowledge, encompassing structural representations (e.g., SMILES strings, molecular graphs) and rich contextual descriptions (e.g., physicochemical properties, biomedical applications) [2, 46, 50]. With rich knowledge integrated, MoLM encoders ingest heterogeneous inputs while decoders generate outputs across modalities, enabling key applications such as molecule-to-caption generation and caption-to-molecule design [34]. However, during knowledge integration, inaccurate or misleading information can be introduced from outdated web-mined training corpora [9], posing risks in downstream applications. For instance, Tox21 [51] initially labeled several toxic molecules as non-toxic based on early assays, which MoLMs may memorize and propagate, potentially leading to critical errors in drug discovery pipelines. Hence, there is a pressing need to refine MoLMs to correct inaccurate or misleading knowledge from outdated training data or malicious manipulation, as shown in Figure 1. Recently, knowledge editing has emerged as an efficient approach to modifying specific knowledge while preserving other information [38, 54, 62]. It has been widely applied to general-domain large language models [7, 40] and multimodal language models [3, 18].

However, the adaptation of knowledge editing to MoLMs remains largely unexplored, which presents two inherent molecule-specific challenges. First, molecular knowledge is inherently *multifaceted* [2]. A single compound consists of multiple functional groups such as aromatic rings, heteroatoms, and chiral centers,



**Figure 1: An illustration of MoLM editing for two tasks: correcting inaccurate captions in molecule-to-caption generation and fixing mismatched or invalid molecules in caption-to-molecule generation.**

each carrying its own reactivity and toxicity signals, while its caption comprises distinct descriptive components such as functional descriptions (e.g., "antiviral diterpenoid") and provenance notes (e.g., "isolated from *Taxus brevifolia*"). Each of these elements represents a specific aspect of molecular properties and may exhibit varying sensitivities to editing. Consequently, modifying multifaceted molecular knowledge poses the risk of over-editing certain aspects while under-editing others [20, 66]. Second, shared functional groups and contextual descriptions create dense *interdependencies* among molecules, so editing one molecule's knowledge can unintentionally affect others with similar features. This makes it difficult to ensure edits remain localized, violating the principle of locality, where modifications should only impact the intended target. For instance, both celecoxib (an anti-inflammatory) and sulfamethoxazole (an antibiotic) share a sulfonamide group, yet they differ markedly in toxicity and side-effect profiles. Editing the model's knowledge of celecoxib might as well alter the model's knowledge of sulfamethoxazole due to their structural similarity, even though the latter should not be edited.

To address these challenges, we propose MolEdit, a simple yet powerful framework for editing multimodal MoLMs<sup>1</sup> in caption generation and molecule design tasks. Our approach enables precise, targeted updates to compositional molecular knowledge while preserving unrelated information. To address the first challenge, we design a *Multi-Expert Knowledge Adapter* (MEKA) that decomposes a molecule or caption into distinct facets and directs different facets of molecular knowledge to specialized editing experts within a lightweight Mixture-of-Expert layer, thus enabling fine-grained control over multifaceted updates. For the second challenge, we introduce an *Expertise-Aware Editing Switcher* (EAES), which maintains a memory bank of edited molecular knowledge and activates the knowledge adapter only when the input closely matches the stored edits across all expertise, thus minimizing interference with

unrelated knowledge. Furthermore, since incorrect outputs can arise from interactions between different modalities, MolEdit edits both encoders and decoders to ensure comprehensive refinement.

To enable systematic evaluation, we introduce MEBench, the first benchmark for editing MoLMs, which rigorously assesses reliability (editing accuracy), locality (preservation of unrelated knowledge), and generality (consistency across varied textual descriptions). Extensive experiments on two popular MoLM backbones show that MolEdit surpasses fine-tuning, gradient-based, and memory-based baselines across all three MEBench dimensions. Overall, our key contributions are as follows:

- **Problem Formulation:** We present the first systematic exploration of *knowledge editing* in molecule language models, clearly delineating two molecule-specific knowledge editing challenges including multifaceted expertise and cross-molecule interdependence.
- **Benchmark Construction:** We release *MEBench*, a rigorously curated dataset that quantifies editing success along three dimensions including *reliability*, *locality*, and *generality*, covering both molecule-to-caption and caption-to-molecule tasks.
- **Framework Design:** We introduce *MolEdit*, a powerful framework that combines a *Multi-Expert Knowledge Adapter* for facet-level editing with an *Expertise-Aware Editing Switcher* that gates updates based on fine-grained similarity, thereby enforcing fine-grained control and locality.
- **Experimental Evaluation:** Through extensive studies on two strong MoLM backbones, MolEdit delivers consistent gains over state-of-the-art gradient-based and external-memory-based baselines on every MEBench metric, demonstrating its effectiveness.

## 2 Related Work

**Molecule Language Models.** Building on the success of general-domain language models, *molecule language models* (MoLMs) aim to learn unified representations that cover structure, property and function [29, 45]. Early single-modal molecule language models such as ChemBERTa and MolT5 concentrated solely on SMILES strings, whereas more recent work shifted into *multimodality* including graph and SMILES encoders that capture molecule structural knowledge and text encoders ingest synthesis notes or pharmacological annotations [30, 47]. To encode multimodal molecular knowledge, two broad pre-training paradigms have emerged. *Generative* schemes treat every modality as a sequence and train an autoregressive decoder that can translate between them, enabling direct sequence-to-sequence tasks [6, 14, 60, 65]. By contrast, *contrastive* schemes retain separate encoders and maximize cross-modal agreement, yielding greater architectural flexibility and lower latency at inference [23, 28, 33–35, 50]. After pre-training, MoLMs are fine-tuned for downstream tasks such as molecule-to-caption generation, caption-to-molecule design, virtual screening, and property prediction [15, 22, 32, 49, 50]. However, both the diverse pre-training corpora and task-specific fine-tuning stages can inject stale or spurious facts such as mis-annotated toxicity labels or incorrect stereochemistry [11, 17]. These issues underscore the

<sup>1</sup>We focus on multimodal MoLMs because they are more challenging and requires edits across both structural and textual representations of molecules.

need for techniques that can edit a model’s knowledge without expensive retraining.

**Knowledge Editing.** A naive way to correct model errors is to *fine-tune* on a small patch dataset, while it can cause costly retraining and catastrophic forgetting of unrelated facts. Knowledge editing therefore focuses on making targeted, data-efficient corrections to a model’s internal parameters while avoiding collateral damage to unrelated facts [38, 61]. There are several types of knowledge editing methods. *Gradient-based* editors adjust weights directly via one or a few back-prop steps, either with meta-learned optimizers [5, 41, 48] or with locate-then-edit strategies that first identify influential sub-spaces [39, 40, 43]. Although effective, they might still distort neighbouring knowledge. *External-memory* approaches instead graft new information onto adapters, side networks, or retrieval memories—reducing interference at the cost of extra parameters [16, 19, 37, 42, 53, 66]. Recent work has begun extending these ideas to multimodal frameworks [4, 59], but none of them are designed to address the *multifaceted* and highly *interdependent* challenge of multimodal chemical knowledge. MolEdit fills this gap with MEKA which routes distinct facets to specialised experts and EAES that activates those experts only when all facets of the query align with a stored edit, thus achieving fine-grained updates while preserving surrounding chemistry.

### 3 Preliminary

**Notations.** Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  represent a molecular graph, where  $\mathcal{V}$  denotes the set of atoms (nodes) and  $\mathcal{E}$  represents covalent bonds (edges). Each molecular graph consists of  $N_g$  subgraphs  $g_{i=1}^{N_g}$ , each corresponding to a functional group. The molecular structure can also be expressed as a SMILES string, denoted as  $\mathcal{S}$ . Additionally, each molecule is associated with a caption  $\mathcal{T} = t^1, \dots, t^{N_t}$  containing  $N_t$  textual descriptions.

**Molecule Language Model.** Given molecular representations  $\mathcal{G}$ ,  $\mathcal{S}$ , and captions  $\mathcal{T}$ , molecule language models learn aligned cross-modal representations by utilizing a molecule and a text encoder during pretraining. A task-specific decoder is then appended for downstream generation tasks.

**Caption Generation.** For molecule-to-caption generation task, a pretrained text decoder is appended to the molecule encoder and fine-tuned to generate captions  $\mathcal{T}$  given a molecular representation  $\mathcal{G}$ ,  $\mathcal{S}$ . We denote the fine-tuned MoLMs for this task as  $f_{cap}$ .

**Molecule Generation.** Similarly, for the caption-to-molecule generation task, a pretrained molecule generation decoder is appended to the text encoder and fine-tuned to generate SMILES representations  $\mathcal{S}$  given a caption  $\mathcal{T}$ <sup>2</sup>. We denote the fine-tuned MoLMs for this task as  $f_{gen}$ .

## 4 Editing Molecule Language Model

In this section, we first introduce the task of editing MoLMs for molecule and caption generation (§4.1). We then present MEBench, the first benchmark for evaluating MoLM editing (§4.2), followed by MolEdit, a simple yet powerful framework designed to address molecule-specific challenges (§4.3).

<sup>2</sup>We follow the definition of existing MoLMs [12, 34, 35]

### 4.1 Task Definition

**Editing Caption Generation.** MoLMs may produce inaccurate outdated captions that require correction. To assess editing effectiveness, we evaluate caption generation along two dimensions. First, **Reliability** measures how well the edited captions align with the expert-curated ground truth. Given a dataset  $\mathcal{D}_{edit}^{\mathcal{T}}$  containing molecules with initially incorrect captions, we compute the semantic similarity between the captions generated by the edited MoLMs,  $\tilde{f}_{cap}$ , and the target captions  $\mathcal{T}$ :

$$\mathcal{M}_{rel}^{\mathcal{T}} = \mathbb{E}_{(\mathcal{G}, \mathcal{S}, \mathcal{T}) \sim \mathcal{D}_{edit}^{\mathcal{T}}} (\mathbf{SIM}_{\mathcal{T}}(\tilde{f}_{cap}(\mathcal{G}, \mathcal{S}), \mathcal{T})), \quad (1)$$

where  $\mathbf{SIM}_{\mathcal{T}}$  is the metric to measure text similarity. Secondly, **Locality** preserves existing knowledge by minimizing the deviations in unedited captions. For a dataset  $\mathcal{D}_{loc}^{\mathcal{T}}$  with knowledge unrelated to  $\mathcal{D}_{edit}^{\mathcal{T}}$ , we compare outputs before and after editing:

$$\mathcal{M}_{loc}^{\mathcal{T}} = \mathbb{E}_{(\mathcal{G}, \mathcal{S}) \sim \mathcal{D}_{loc}^{\mathcal{T}}} (\mathbf{SIM}_{\mathcal{T}}(\tilde{f}_{cap}(\mathcal{G}, \mathcal{S}), f_{cap}(\mathcal{G}, \mathcal{S}))). \quad (2)$$

**Editing Molecule Generation.** MoLMs can also generate invalid molecule SMILES that require correction. Similar to editing caption generation, we evaluate the **Reliability** and **Locality** with edited MoLMs for molecule generation  $\tilde{f}_{gen}$ :

$$\mathcal{M}_{rel}^{\mathcal{S}} = \mathbb{E}_{(\mathcal{S}, \mathcal{T}) \sim \mathcal{D}_{edit}^{\mathcal{S}}} (\mathbf{SIM}_{\mathcal{G}}(\tilde{f}_{gen}(\mathcal{T}), \mathcal{S})), \quad (3)$$

$$\mathcal{M}_{loc}^{\mathcal{S}} = \mathbb{E}_{(\mathcal{T}) \sim \mathcal{D}_{loc}^{\mathcal{S}}} (\mathbf{SIM}_{\mathcal{G}}(\tilde{f}_{gen}(\mathcal{T}), f_{gen}(\mathcal{T}))), \quad (4)$$

where  $\mathcal{D}_{edit}^{\mathcal{S}}$  contain molecules requiring knowledge editing and  $\mathcal{D}_{loc}^{\mathcal{S}}$  ought to remain unchanged during editing.  $\mathbf{SIM}_{\mathcal{G}}$  is the metric to measure molecule similarity. Additionally, to prevent the model from merely memorizing descriptions during training, semantically equivalent descriptions should be designed to generate the same molecule. To evaluate this, we assess the model’s output consistency using a **generality** dataset, as shown below:

$$\mathcal{M}_{gen}^{\mathcal{S}} = \mathbb{E}_{(\mathcal{T}_r) \sim \mathcal{N}(\mathcal{T})} (\mathbf{SIM}_{\mathcal{G}}(\tilde{f}_{gen}(\mathcal{T}_r), \mathcal{S})), \quad (5)$$

( $\mathcal{S}, \mathcal{T}$ )  $\sim \mathcal{D}_{edit}^{\mathcal{S}}$

where the  $\mathcal{N}(\cdot)$  denotes the description generalization set.

### 4.2 Benchmark Construction

This subsection provides a brief overview of the MEBench construction procedure. For clarity, Figure 2 depicts a representative sample drawn from the final benchmark.

#### 4.2.1 Editing Caption Generation.

**Reliability.** To assess the effectiveness of knowledge editing, we construct a caption reliability dataset,  $\mathcal{D}_{edit}^{\mathcal{T}}$ . We first identify suboptimal entries (BLEU-2 score [55]  $< 0.2$ ) shared across multiple MoLMs within the widely used molecule dataset (ChEBI-20 [12]) for caption generation, using its ground truth captions as the editing targets. Additionally, to enable a more fine-grained evaluation of editing capabilities (as discussed in Section 5.4), the targets are decomposed into distinct descriptions, each capturing a specific aspect of molecular knowledge.

| Reliability                                                                                                                                                                                                                                      | Reliability                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input:</b> <chem>CN(C)CCCN1C2=CC=CC=C2C3=C1C=C(C=C3)Cl</chem><br><b>Target:</b> The molecule is a dibenzoazepine. One of the more sedating tricyclic antidepressants. It derives from an imipramine. It is a conjugate of a clomipramine(1+). | <b>Input:</b> The molecule is a member of the class of thioureas. It is a pyrimidinecarboxylate ester.<br><b>Target:</b> <chem>CCOC(=O)C1=C(NC(=S)NC1C2=CC(=CC=C2)OC)C</chem>       |
| Locality                                                                                                                                                                                                                                         | Semantic Generality                                                                                                                                                                 |
| <b>Input:</b> <chem>CC(CN1C2=CC=CC=C2SC3=CC=CC=C31)N(C)C</chem><br><b>Target:</b> This molecule is an ammonium ion derivative and an organic cation. It is a conjugate acid of a clomipramine.                                                   | <b>Input:</b> The molecule belongs to thioureas. It has an ester attached to a pyrimidinecarboxylate backbone<br><b>Target:</b> <chem>CCOC(=O)C1=C(NC(=S)NC1C2=CC(=CC=C2)O)C</chem> |
| Editing Caption Generation                                                                                                                                                                                                                       | Editing Molecule Generation                                                                                                                                                         |

**Figure 2: A sample illustration of MEBench. It includes three evaluation dimensions for two tasks: Reliability (molecules requiring editing), Locality (similar but untargeted molecules), and Generality (semantically equivalent but different captions).**

**Locality.** To assess the ability of MoLMs to preserve existing knowledge, we construct a locality dataset,  $\mathcal{D}_{\text{loc}}^T$ . Specifically, we extract high-accuracy entries (BLEU-2 score > 0.95) from the training set to serve as the basis for this dataset. Since knowledge editing is more likely to affect knowledge similar to the editing target, we select locality samples with high similarity to those in the reliability dataset, making the editing task more challenging. Following the design of UnKEBench [8], both the reliability and locality datasets contain unstructured knowledge, as molecule captions are complex and involve multiple entities.

#### 4.2.2 Editing Molecule Generation.

**Reliability.** To assess the effectiveness of knowledge editing in improving molecule SMILES generation, we construct the molecule reliability dataset,  $\mathcal{D}_{\text{edit}}^S$ . Following a similar approach to caption editing dataset construction, we identify suboptimal entries (MACCS FTS [64] < 0.2) shared across multiple MoLMs within the same dataset. The ground truth SMILES representations in the dataset serve as the editing targets.

**Locality.** To construct the generation locality dataset,  $\mathcal{D}_{\text{loc}}^S$ , we select high-accuracy entries (MACCS FTS > 0.95) from multiple MoLMs within the training set. To ensure a rigorous evaluation of the model’s ability to preserve existing knowledge, we choose entries with the highest similarity to those in the reliability dataset.

**Generality.** To assess a model’s ability to generalize across different phrasings, we construct the generality dataset,  $\mathcal{N}(\mathcal{T})$ . This dataset consists of rephrased captions from the molecule reliability dataset, evaluating whether semantically equivalent descriptions consistently generate the same molecular structures.

### 4.3 Methodology

To tackle the challenges of editing MoLMs, we propose MolEdit, a simple yet powerful framework designed to enable localized updates to multifaceted molecular knowledge. As shown in Figure 3,

MolEdit comprises two key components: (1) *Multi-Expert Knowledge Adapter* - This module disentangles and customizes the editing of diverse molecular knowledge (e.g., functional groups in molecules, descriptive elements in captions) by dynamically routing them to specialized editing experts via a Mixture-of-Experts (MoE) architecture. (2) *Expertise-Aware Editing Switcher* - This component ensures edits are applied only to highly relevant inputs by leveraging a memory bank of expertise embeddings, activating modifications only when the input exhibits substantial overlap with stored edits.

**4.3.1 Multi-Expert Knowledge Adapter.** Since errors can originate from both input and output modalities [3], MolEdit enables knowledge editing by wrapping selected layers in both the encoder and decoder with adapters, allowing localized parameter updates. Taking the molecule generation editing task as an example, when editing the encoder at layer  $l$ , we employ  $P$  distinct experts to perform customized edits for different descriptions in the input. Each expert is instantiated as a feed-forward network (FFN) layer. To achieve this, a gating function dynamically routes the  $(l-1)$ -th layer embeddings of each description to the appropriate expert, illustrated as,

$$\mathcal{P}^n = \text{top}_k \left( \text{softmax} \left( \mathbf{W}_g \cdot \Sigma_{i \in t_n} (z_i^{l-1}) / |t^n| + \epsilon \right) \right), \quad (6)$$

where  $t^n$  is the  $n$ -th description and  $z_i^{l-1}$  is the embedding of  $i$ -th token in  $t^n$  at  $(l-1)$ -th layer.  $\mathbf{W}_g$  is the trainable weights in gate decision, while  $\epsilon$  denotes the noise term. The  $\text{top}_k(\cdot)$  operator zeros out all but the top- $k$  values. After getting the gate decision vector  $\mathcal{G}^n$  of the  $n$ -th description, the corresponding output is generated through a weighted aggregation of each expert’s computation on  $z_i^l$ , demonstrated as,


$$z_i^l = f^l(z_i^{l-1}) + \sum_{p=1}^P \mathcal{P}_p^n \cdot \mathbf{W}_p \cdot z_i^{l-1}, \quad i \in t^n, \quad (7)$$

where  $\mathbf{W}_p$  is the trainable weights. When editing the decoder at layer  $l'$ , the ground truth expertise segmentation is unavailable. As a result, the MoE adapter is applied to each token, which is expected to route tokens associated with different expertise to the appropriate experts, demonstrated as,

$$\begin{aligned} \mathcal{P}^i &= \text{top}_k \left( \text{softmax} \left( \mathbf{W}_g' \cdot z_i^{l'-1} + \epsilon \right) \right), \\ z_i^{l'} &= f^{l'}(z_i^{l'-1}) + \lambda \sum_{p=1}^P \mathcal{P}_p^i \cdot \mathbf{W}_p' \cdot z_i^{l'-1}. \end{aligned} \quad (8)$$

Similarly, for caption generation, the input molecule is composed of multiple functional groups, each representing a distinct molecule expertise. We route by functional groups during encoder editing and by token during decoder editing due to the lack of predefined expertise in decoder.

**4.3.2 Expertise-Aware Editing Switcher.** In order to minimize unintended interference with untargeted molecules that share a few functional groups or descriptions, we design an expertise-aware switching mechanism that only allows activation of the knowledge adapters for molecules with high expertise overlap with edited molecules. Specifically, we keep an expertise-based memory bank



**Figure 3: An overview of MolEdit to edit MoLMs for molecule/caption generation by modifying a chosen layer in either the encoder and decoder. It is composed of two components: (1) Multi-Expert Knowledge Adapter (MEKA) and (2) Expertise-Aware Editing Switcher (EAES). Specifically, MEKA utilizes expertise-wise MoE for encoder and token-wise MoE for decoder to route expertise to different editing experts (instantiated as FFN). EAES stores edited knowledge expertise (functional groups/descriptions) and activates MEKA only when all input expertise finds a similar match in its memory bank during inference, thereby preserving unrelated knowledge.**

that stores the expertise-wise knowledge through the form of encoder embeddings of the edit queries:

$$\mathcal{Z} = \{\bar{z}_{n_j}\}, \quad \bar{z}_{n_j} = \sum_{i \in t_{n_j}} (z_i^{enc}) / |t_{n_j}|, \quad (9)$$

where  $z_i^{enc}$  represents the encoder embedding of the  $i$ -th token (node) within the  $n_j$ -th description (functional group) of the  $j$ -th sample. During inference, the switcher compares each expertise in the input to the expertise of stored edits in the memory bank  $\mathcal{Z}$ . The adapter is activated only when all expertise distances are below a threshold  $\epsilon$ , ensuring that unrelated molecules remain unaffected:

$$z_i^l = \begin{cases} \text{MolEdit}(z_i^{l-1}) & \text{if } \max_n (d(\bar{z}_n, \mathcal{Z})) < \epsilon, \\ f^l(z_i^{l-1}) & \text{otherwise,} \end{cases} \quad (10)$$

where  $d(\cdot)$  is a distance function.

## 5 Experiments

We aim to answer three key research questions: **RQ1**: How does MolEdit perform compared to baselines on MEBench? **RQ2**: How does each component contribute to MolEdit’s performance? **RQ3**: How can we explain the effectiveness of each module? In addition, we also include qualitative generation examples contrasting MolEdit with baseline outputs. The analysis is presented below.

### 5.1 Experiment Settings

Below we introduce the experiment settings, including MoLM backbones, baselines, metrics, and implementation details.

**5.1.1 MoLM Backbone.** In this paper, we adopt two widely utilized multimodal MoLMs with different paradigms as the editing backbones: MoMu [50] and MolFM [34].

**MoMu.** MoMu is a multimodal MoLM aligning molecule graphs and text through contrastive learning [24]. MoMu utilizes a GIN encoder for graphs [57] and a Bert encoder for text [10]. For generation, MoMu attaches a MolT5 decoder [12], enabling both molecule-to-caption and caption-to-molecule tasks.

**MolFM.** MolFM integrates molecular structures, biomedical text, and knowledge graphs using cross-modal attention to fuse signals across modalities. It uses GraphMVP as the graph encoder [31] and BERT for text, with MolT5 as a unified decoder for caption and molecule generation. Compared with MoMu’s contrastive alignment, MolFM’s attention-based fusion yields richer token-level interactions during pretraining, offering a complementary testbed for molecular knowledge editing.

**5.1.2 Baselines. Fine-tune.** Fine-tune adapts pre-trained language models to specific tasks and remains the default knowledge editing baseline [25, 36, 67]. We experiment with three fine-tuning variants: (1) *encoder-only*, (2) *decoder-only*, and (3) *encoder+decoder* editing.

**MEND.** MEND [41] is an *optimization-based method* that adopts a mechanism referred to as gradient decomposition. It employs auxiliary editing networks that transform model gradients via low-rank decomposition. We adapt MEND to our multimodal setting by applying edits at both encoder and decoder layers.

**GRACE.** GRACE [16] is a *memory-based method* that employs a deferral mechanism to decide whether to activate an adapter by comparing a given input against a codebook of stored edits. Similar to MEND, we extend GRACE to the multi-modal setting.

**5.1.3 Metrics.** We evaluate editing quality on MEBench using standard MoLM metrics, grouped by task. For editing molecule generation, we report BLEU-4 [44] to measure string-level corrections, LEV [58] to capture normalized edit distance that calculates

|     |                | MoMu                         |                               |                              | MolFM                        |                               |                              |
|-----|----------------|------------------------------|-------------------------------|------------------------------|------------------------------|-------------------------------|------------------------------|
|     |                | BLEU-4 $\uparrow$            | LEV $\downarrow$              | MACCS $\uparrow$             | BLEU-4 $\uparrow$            | LEV $\downarrow$              | MACCS $\uparrow$             |
| Rel | FT (Encoder)   | 0.758 ( $\pm 0.022$ )        | 22.974 ( $\pm 3.224$ )        | 0.987 ( $\pm 0.018$ )        | <b>0.983</b> ( $\pm 0.000$ ) | <b>2.015</b> ( $\pm 0.018$ )  | <u>0.998</u> ( $\pm 0.001$ ) |
|     | FT (Decoder)   | 0.711 ( $\pm 0.002$ )        | 29.936 ( $\pm 0.112$ )        | 0.938 ( $\pm 0.002$ )        | 0.894 ( $\pm 0.035$ )        | 11.560 ( $\pm 4.244$ )        | 0.977 ( $\pm 0.013$ )        |
|     | FT (All)       | 0.781 ( $\pm 0.000$ )        | <u>19.940</u> ( $\pm 0.035$ ) | <b>1.000</b> ( $\pm 0.000$ ) | <u>0.977</u> ( $\pm 0.011$ ) | <u>2.554</u> ( $\pm 1.091$ )  | 0.995 ( $\pm 0.001$ )        |
|     | MEND           | <u>0.802</u> ( $\pm 0.019$ ) | 21.079 ( $\pm 1.360$ )        | 0.834 ( $\pm 0.011$ )        | 0.789 ( $\pm 0.003$ )        | 23.547 ( $\pm 0.125$ )        | 0.869 ( $\pm 0.003$ )        |
|     | GRACE          | 0.718 ( $\pm 0.000$ )        | 28.331 ( $\pm 0.025$ )        | 0.938 ( $\pm 0.000$ )        | 0.770 ( $\pm 0.001$ )        | 11.464 ( $\pm 15.122$ )       | 0.987 ( $\pm 0.004$ )        |
|     | <b>MolEdit</b> | <b>0.953</b> ( $\pm 0.025$ ) | <b>4.667</b> ( $\pm 2.154$ )  | <u>0.989</u> ( $\pm 0.008$ ) | 0.975 ( $\pm 0.003$ )        | 2.862 ( $\pm 0.479$ )         | <b>1.000</b> ( $\pm 0.000$ ) |
| Loc | FT (Encoder)   | 0.829 ( $\pm 0.001$ )        | 18.786 ( $\pm 0.089$ )        | 0.881 ( $\pm 0.008$ )        | 0.829 ( $\pm 0.001$ )        | 19.622 ( $\pm 0.154$ )        | <u>0.943</u> ( $\pm 0.005$ ) |
|     | FT (Decoder)   | 0.881 ( $\pm 0.001$ )        | 12.562 ( $\pm 0.094$ )        | 0.936 ( $\pm 0.003$ )        | 0.725 ( $\pm 0.001$ )        | 31.195 ( $\pm 0.203$ )        | 0.826 ( $\pm 0.002$ )        |
|     | FT (All)       | 0.891 ( $\pm 0.004$ )        | 11.756 ( $\pm 0.413$ )        | <u>0.949</u> ( $\pm 0.005$ ) | 0.803 ( $\pm 0.009$ )        | 21.985 ( $\pm 1.189$ )        | 0.909 ( $\pm 0.006$ )        |
|     | MEND           | <u>0.912</u> ( $\pm 0.028$ ) | <u>9.512</u> ( $\pm 3.291$ )  | 0.940 ( $\pm 0.030$ )        | <u>0.833</u> ( $\pm 0.020$ ) | <u>17.581</u> ( $\pm 1.750$ ) | 0.937 ( $\pm 0.014$ )        |
|     | GRACE          | 0.859 ( $\pm 0.000$ )        | 15.732 ( $\pm 0.005$ )        | 0.937 ( $\pm 0.000$ )        | 0.757 ( $\pm 0.001$ )        | 30.112 ( $\pm 0.121$ )        | 0.894 ( $\pm 0.015$ )        |
|     | <b>MolEdit</b> | <b>0.980</b> ( $\pm 0.007$ ) | <b>1.853</b> ( $\pm 0.727$ )  | <b>0.997</b> ( $\pm 0.000$ ) | <b>0.918</b> ( $\pm 0.034$ ) | <b>9.167</b> ( $\pm 3.963$ )  | <b>0.991</b> ( $\pm 0.007$ ) |
| Gen | FT (Encoder)   | 0.526 ( $\pm 0.013$ )        | 55.008 ( $\pm 0.359$ )        | 0.629 ( $\pm 0.012$ )        | 0.727 ( $\pm 0.021$ )        | <u>30.209</u> ( $\pm 1.162$ ) | 0.697 ( $\pm 0.051$ )        |
|     | FT (Decoder)   | 0.590 ( $\pm 0.011$ )        | 48.795 ( $\pm 1.465$ )        | 0.753 ( $\pm 0.021$ )        | 0.665 ( $\pm 0.002$ )        | 38.035 ( $\pm 0.172$ )        | 0.649 ( $\pm 0.023$ )        |
|     | FT (All)       | 0.586 ( $\pm 0.006$ )        | 47.872 ( $\pm 1.192$ )        | 0.745 ( $\pm 0.013$ )        | 0.713 ( $\pm 0.002$ )        | 30.755 ( $\pm 2.094$ )        | 0.691 ( $\pm 0.017$ )        |
|     | MEND           | <u>0.707</u> ( $\pm 0.025$ ) | <u>30.551</u> ( $\pm 2.379$ ) | 0.721 ( $\pm 0.009$ )        | <u>0.731</u> ( $\pm 0.000$ ) | 30.945 ( $\pm 0.513$ )        | 0.678 ( $\pm 0.052$ )        |
|     | GRACE          | 0.703 ( $\pm 0.000$ )        | 32.574 ( $\pm 0.000$ )        | <u>0.913</u> ( $\pm 0.000$ ) | 0.645 ( $\pm 0.000$ )        | 44.521 ( $\pm 0.732$ )        | <u>0.892</u> ( $\pm 0.029$ ) |
|     | <b>MolEdit</b> | <b>0.842</b> ( $\pm 0.028$ ) | <b>15.609</b> ( $\pm 2.304$ ) | <b>0.917</b> ( $\pm 0.020$ ) | <b>0.796</b> ( $\pm 0.020$ ) | <b>23.314</b> ( $\pm 1.576$ ) | <b>0.895</b> ( $\pm 0.044$ ) |

**Table 1: Main results on MEBench for editing MoMu and MolFM in molecule generation, evaluated across three dimensions: Reliability (Rel), Locality (Loc), and Generality (Gen). Each dimension uses three metrics: BLEU-4, LEV, and MACSS. The best and second-best results are shown in bold and underlined, respectively. FT denotes fine-tuning.**

the length-normalized Levenshtein distance between two SMILES strings, and MACCS FTS [64] to assess Tanimoto substructure similarity via MACCS fingerprint matches. Together, these metrics provide complementary lexical and structural perspectives on the quality of the edits. For editing caption generation, we employ BLEU-2[55] to measure n-gram overlap, METEOR[1] to account for semantic similarity through synonym and stem matching, and ROUGE-1 [27] to evaluate unigram recall, offering a balanced assessment of lexical and semantic fidelity.

**5.1.4 Implementation Details.** In this section, we will introduce the implementation details of MolEdit. The Multi-Expert Knowledge Adapter’s similarity threshold  $\epsilon$  is 0.9 for MolFM molecule generation, 0.8 for MoMu molecule generation, and 0.9 for caption generation. The number of experts  $P$  is set to 5, with top- $k$  fixed at 1. The distance function  $d(\cdot)$  is instantiated as cosine similarity. Following prior work [3, 52], we edit mid-to-late layers (specifically, layer 4 of the encoder and layer 10 of the decoder in our work) for all tasks and backbones. We edit one piece of knowledge at a time for molecule generation and two pieces for caption generation, due to batch normalization in the MoMu and MolFM molecule encoders. In addition, learning rates are set to  $1e-4$  for caption generation,  $2e-5$  for MoMu molecule generation, and  $1e-5$  for MolFM molecule generation. All experiments were conducted on an NVIDIA A100 server with four GPUs.

## 5.2 Main Experiments

To answer **RQ1**, we first evaluate MolEdit’s performance against all baseline methods on MEBench for both molecule and caption

generation tasks. We make the following observations in Table 1 and Table 2: (1) For molecule generation editing, MolEdit consistently surpasses all baselines across most metrics. In particular, under the BLEU-4 metric, MolEdit outperforms the second-best approach by up to 18.8% in Reliability, 10.2% in Locality, and 19.1% in Generality. (2) For caption generation editing, MolEdit again achieves superior performance across all metrics. Notably, in terms of BLEU-2, it exceeds the second-best method by up to 5.4% in Reliability and 12.0% in Locality. These results highlight MolEdit’s effectiveness in accurately updating molecule-specific knowledge across different editing tasks. (3) MolEdit generally performs worse on molecule generation than on caption generation, possibly because editing molecular SMILES strings is inherently more complex and abstract, lacking explicit semantic structure. (4) MolEdit performs better when editing MoMu compared to MolFM, suggesting that editing becomes more challenging as the underlying model complexity increases. (5) The performance under the Generality is consistently worse than Reliability, indicating a persistent gap across different representations that encode the same underlying knowledge. We further observe differences in how baselines behave on MEBench: (1) Different methods excel in different evaluation dimensions on MEBench. Fine-tuning, while effective at reliable knowledge updates, struggles to preserve irrelevant knowledge, since it updates a large portion of the model’s parameters without constraint. MEND performs better at knowledge preservation but underperforms in Reliability, potentially due to the complexity of unstructured molecular knowledge, which poses a challenge for meta-learning approaches. GRACE excels at editing caption

|     |                | MoMu                         |                              |                              | MolFM                        |                              |                              |
|-----|----------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|
|     |                | BLEU-2↑                      | METEOR↑                      | ROUGE-1↑                     | BLEU-2↑                      | METEOR↑                      | ROUGE-1↑                     |
| Rel | FT (Encoder)   | 0.334 ( $\pm 0.005$ )        | 0.365 ( $\pm 0.006$ )        | 0.472 ( $\pm 0.007$ )        | 0.321 ( $\pm 0.014$ )        | 0.346 ( $\pm 0.019$ )        | 0.448 ( $\pm 0.023$ )        |
|     | FT (Decoder)   | 0.784 ( $\pm 0.017$ )        | 0.813 ( $\pm 0.014$ )        | 0.854 ( $\pm 0.013$ )        | 0.916 ( $\pm 0.000$ )        | 0.937 ( $\pm 0.000$ )        | 0.958 ( $\pm 0.000$ )        |
|     | FT (All)       | 0.886 ( $\pm 0.034$ )        | 0.907 ( $\pm 0.030$ )        | 0.925 ( $\pm 0.024$ )        | 0.921 ( $\pm 0.001$ )        | 0.941 ( $\pm 0.001$ )        | 0.960 ( $\pm 0.001$ )        |
|     | MEND           | 0.557 ( $\pm 0.034$ )        | 0.569 ( $\pm 0.024$ )        | 0.631 ( $\pm 0.022$ )        | 0.627 ( $\pm 0.025$ )        | 0.652 ( $\pm 0.014$ )        | 0.715 ( $\pm 0.007$ )        |
|     | GRACE          | <u>0.928</u> ( $\pm 0.000$ ) | <u>0.947</u> ( $\pm 0.000$ ) | <u>0.965</u> ( $\pm 0.000$ ) | <u>0.928</u> ( $\pm 0.000$ ) | <u>0.947</u> ( $\pm 0.000$ ) | <u>0.965</u> ( $\pm 0.000$ ) |
|     | <b>MolEdit</b> | <b>0.978</b> ( $\pm 0.006$ ) | <b>0.978</b> ( $\pm 0.007$ ) | <b>0.982</b> ( $\pm 0.005$ ) | <b>0.977</b> ( $\pm 0.006$ ) | <b>0.979</b> ( $\pm 0.004$ ) | <b>0.983</b> ( $\pm 0.004$ ) |
|     |                |                              |                              |                              |                              |                              |                              |
| Loc | FT (Encoder)   | 0.511 ( $\pm 0.006$ )        | 0.536 ( $\pm 0.010$ )        | 0.604 ( $\pm 0.006$ )        | 0.491 ( $\pm 0.054$ )        | 0.508 ( $\pm 0.065$ )        | 0.587 ( $\pm 0.052$ )        |
|     | FT (Decoder)   | 0.637 ( $\pm 0.023$ )        | 0.653 ( $\pm 0.030$ )        | 0.699 ( $\pm 0.023$ )        | 0.669 ( $\pm 0.000$ )        | 0.688 ( $\pm 0.000$ )        | 0.732 ( $\pm 0.000$ )        |
|     | FT (All)       | 0.625 ( $\pm 0.066$ )        | 0.637 ( $\pm 0.063$ )        | 0.688 ( $\pm 0.056$ )        | 0.656 ( $\pm 0.004$ )        | 0.671 ( $\pm 0.006$ )        | 0.721 ( $\pm 0.007$ )        |
|     | MEND           | 0.615 ( $\pm 0.010$ )        | 0.633 ( $\pm 0.011$ )        | 0.676 ( $\pm 0.016$ )        | 0.844 ( $\pm 0.005$ )        | 0.856 ( $\pm 0.006$ )        | 0.873 ( $\pm 0.005$ )        |
|     | GRACE          | <u>0.875</u> ( $\pm 0.013$ ) | <u>0.883</u> ( $\pm 0.006$ ) | <u>0.904</u> ( $\pm 0.005$ ) | <u>0.893</u> ( $\pm 0.002$ ) | <u>0.899</u> ( $\pm 0.001$ ) | <u>0.917</u> ( $\pm 0.002$ ) |
|     | <b>MolEdit</b> | <b>0.978</b> ( $\pm 0.010$ ) | <b>0.981</b> ( $\pm 0.009$ ) | <b>0.983</b> ( $\pm 0.008$ ) | <b>0.991</b> ( $\pm 0.001$ ) | <b>0.991</b> ( $\pm 0.000$ ) | <b>0.993</b> ( $\pm 0.000$ ) |
|     |                |                              |                              |                              |                              |                              |                              |

Table 2: Main results on MEBench for editing MoMu and MolFM in caption generation, evaluated across Reliability (Rel) and Locality (Loc) dimensions. Each dimension uses three metrics: BLEU-2, METEOR, and ROUGE-1. The best and second-best results are shown in bold and underlined, respectively. FT denotes fine-tuning.

|     |              | MoMu                         |                              |                              | MolFM                        |                              |                              |
|-----|--------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|
|     |              | BLEU-2↑                      | METEOR↑                      | ROUGE-1↑                     | BLEU-2↑                      | METEOR↑                      | ROUGE-1↑                     |
| Rel | MolEdit      | <b>0.978</b> ( $\pm 0.006$ ) | <b>0.978</b> ( $\pm 0.007$ ) | <b>0.982</b> ( $\pm 0.005$ ) | <b>0.977</b> ( $\pm 0.006$ ) | <b>0.979</b> ( $\pm 0.004$ ) | <b>0.983</b> ( $\pm 0.004$ ) |
|     | w/o EAES     | <u>0.966</u> ( $\pm 0.007$ ) | <u>0.968</u> ( $\pm 0.006$ ) | <u>0.974</u> ( $\pm 0.006$ ) | 0.969 ( $\pm 0.005$ )        | 0.970 ( $\pm 0.005$ )        | 0.975 ( $\pm 0.005$ )        |
|     | w/o MEKA     | 0.773 ( $\pm 0.017$ )        | 0.782 ( $\pm 0.015$ )        | 0.823 ( $\pm 0.013$ )        | 0.803 ( $\pm 0.018$ )        | 0.818 ( $\pm 0.017$ )        | 0.854 ( $\pm 0.015$ )        |
|     | Decoder-only | 0.955 ( $\pm 0.012$ )        | 0.953 ( $\pm 0.011$ )        | 0.964 ( $\pm 0.010$ )        | <u>0.972</u> ( $\pm 0.004$ ) | <u>0.974</u> ( $\pm 0.004$ ) | <u>0.982</u> ( $\pm 0.004$ ) |
|     | Encoder-only | 0.404 ( $\pm 0.014$ )        | 0.416 ( $\pm 0.015$ )        | 0.512 ( $\pm 0.016$ )        | 0.413 ( $\pm 0.016$ )        | 0.426 ( $\pm 0.017$ )        | 0.521 ( $\pm 0.018$ )        |
| Loc | MolEdit      | <b>0.978</b> ( $\pm 0.010$ ) | <b>0.981</b> ( $\pm 0.009$ ) | <b>0.983</b> ( $\pm 0.008$ ) | <u>0.991</u> ( $\pm 0.001$ ) | <u>0.991</u> ( $\pm 0.000$ ) | <u>0.993</u> ( $\pm 0.000$ ) |
|     | w/o EAES     | 0.924 ( $\pm 0.013$ )        | 0.928 ( $\pm 0.012$ )        | 0.944 ( $\pm 0.011$ )        | 0.973 ( $\pm 0.004$ )        | 0.974 ( $\pm 0.004$ )        | 0.977 ( $\pm 0.004$ )        |
|     | w/o MEKA     | <b>0.978</b> ( $\pm 0.008$ ) | <u>0.979</u> ( $\pm 0.007$ ) | <u>0.981</u> ( $\pm 0.007$ ) | 0.990 ( $\pm 0.001$ )        | 0.988 ( $\pm 0.001$ )        | 0.992 ( $\pm 0.001$ )        |
|     | Decoder-only | <u>0.958</u> ( $\pm 0.010$ ) | 0.961 ( $\pm 0.010$ )        | 0.962 ( $\pm 0.009$ )        | 0.977 ( $\pm 0.004$ )        | 0.979 ( $\pm 0.004$ )        | 0.981 ( $\pm 0.004$ )        |
|     | Encoder-only | 0.947 ( $\pm 0.012$ )        | 0.951 ( $\pm 0.011$ )        | 0.952 ( $\pm 0.011$ )        | <b>0.996</b> ( $\pm 0.001$ ) | <b>0.996</b> ( $\pm 0.000$ ) | <b>0.998</b> ( $\pm 0.000$ ) |

Table 3: Ablation study for editing caption generation under the Reliability (Rel) and Locality (Loc) dimensions. For each dimension, we perform the evaluation by using three metrics: BLEU-2, METEOR, and ROUGE-1. The best and second-best results are shown in bold and underlined, respectively.

generation, but performs poorly in editing molecule generation, likely because it treats captions in a coarse-grained manner rather than capturing fine-grained expertise-level distinctions. (2) The choice of the edited module significantly influences performance. Fine-tuning the encoder generally enhances performance in editing molecule generation but leads to weaker performance in editing caption generation, suggesting that knowledge is stored in different locations depending on the tasks. Furthermore, the superior performance achieved by editing both modules highlights the presence of synergistic knowledge storage across them.

### 5.3 Ablation Study

To answer **RQ2**, we assess the contribution of each module in MolEdit to the performance. We use *w/o MEKA* to denote the use of a plain LoRA adapter without MoE and *w/o EAES* to denote

calculating similarity at the sample level rather than the expertise level. Additionally, we evaluate the impact of editing only the encoder or decoder in MolEdit. The results for molecule generation are presented in Figure 4 and the results for caption generation are presented in Table 3. We make the following observations: (1) Both MEKA and EAES contribute to overall performance, validating their effectiveness in addressing molecule-specific challenges in MoLM editing. (2) EAES enhances Locality performance, indicating that it successfully defers untargeted inputs to unedited layers. MEKA improves Reliability performance, suggesting that multiple experts are better equipped to handle multifaceted molecular knowledge. (3) Editing both the encoder and decoder outperforms editing either component individually, highlighting synergistic knowledge storage across these components. (4) Editing the encoder yields markedly lower reliability than editing the decoder, especially for




Figure 4: Ablation study for editing molecule generation under the Reliability, Locality, and Generality dimensions. For each dimension, we perform the evaluation by using three metrics: BLEU-4, LEV, and MACCS. EAES denotes Expertise-Aware Editing Switcher while MEKA denotes Multi-Expert Knowledge Adapter.




Figure 5: Performance of fine-tuning MoMu on a variation dataset of MEBench. Each subset in this variation dataset targets caption editing where only a single type of expertise requires modification. The expertise is labeled by domain experts and describes different molecular aspects, including (1) Function (Func), (2) Origin (Orig), (3) Structure (Stru), (4) Type, and (5) Property (Prop).

caption generation, suggesting that task-relevant knowledge is concentrated on the decoder side.

#### 5.4 Rationale Validation

To answer RQ3, we aim to provide a deeper analysis of each module's rationale. We first validate two interconnected principles




Figure 6: The activation distribution of the experts under different settings. For the label in x-axis, "M" denotes MoMu, "F" denotes MolFM; "C" denotes editing caption generation, "S" denotes editing molecule generation; "E" denotes editing encoder, "D" denotes decoder.

behind the multi-expert knowledge adapter: (1) **The necessity of expertise-wise editing.** We hypothesize that distinct expertise types exhibit heterogeneous sensitivities to editing, creating asymmetric risks of over- and under-editing, thereby necessitating expertise-specific adjustments. To validate this, we curate a specialized dataset from MEBench for molecular caption generation, where each subset focuses on modifying a single targeted expertise,


| Description                                                                                                                                                                                                                                                                                                                                                                                                                         | GRACE                                                                                                                                                                                                                                                                                                                                                    | MolEdit                                                                                                                                                                                                                                                                           | Ground Truth                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The molecule is a terpene lactone—specifically 13,28-epoxyolean-11-en-28-one bearing hydroxy groups at positions 3 and 23—that, having been isolated from the leaves and twigs of <i>Fatsia polycarpa</i> , demonstrates anti-HBV activity, serves as both a general and plant metabolite as well as an anti-HBV agent, is classifiable as a diol and a hexacyclic triterpenoid, and ultimately derives from a hydride of oleanane. | <chem>C[C@]12CC[C@H]([C@@]([C@@H]1CC[C@]3([C@@H]2C=C[C@]45[C@]4C[C@]6([C@H]4CC([C@]4)C([C@]([C@@])C[C@]43[C@]5(C[C@H]([C@@H]4O) adopt CC[C@]CN5)</chem>                                                                                                                                                                                                  | <chem>C[C@]12CC[C@H]([C@@]([C@@H]1CC[C@]3([C@@H]2C=C[C@]45[C@]3(CC[C@]6([C@H]4CC(CCC6)(C)C(=O)O5)C)C)(C)C[C@H]4O)</chem>                                                                                                                                                          | <chem>C[C@]12CC[C@H]([C@@]([C@@H]1CC[C@]3([C@@H]2C=C[C@]45[C@]3(CC[C@]6([C@H]4CC(CCC6)(C)C(=O)O5)C)C(CO)O</chem>                                                                                                                                                                   |
| Molecule                                                                                                                                                                                                                                                                                                                                                                                                                            | GRACE                                                                                                                                                                                                                                                                                                                                                    | MolEdit                                                                                                                                                                                                                                                                           | Ground Truth                                                                                                                                                                                                                                                                       |
|  <p>CCCCCCCCC/C=C/<br/>CCCCCCCC(=O)[O-]</p>                                                                                                                                                                                                                                                                                                        | The molecule is an unsaturated fatty acid anion that is the conjugate base of gadelaic acid, arising from deprotonation of the carboxy group; <b>major species at pH 7.3</b> . It is an unsaturated fatty acid anion, <b>an enol</b> , <b>an aliphatic fatty acid</b> anion and a <b>gadelaicidic acid</b> anion. It is a conjugate base of a gadelaicid | The molecule is an unsaturated fatty acid anion that is the conjugate base of gadelaic acid, arising from deprotonation of the carboxy group. It is an unsaturated fatty acid anion, an icosanoid anion and a long-chain fatty acid anion. It is a conjugate base of a gadelaicid | The molecule is an unsaturated fatty acid anion that is the conjugate base of gadelaic acid, arising from deprotonation of the carboxy group. It is an unsaturated fatty acid anion, an icosanoid anion and a long-chain fatty acid anion. It is a conjugate base of a gadelaicid. |

Figure 7: Case study of two tasks on MolFM. Errors are marked in red. GRACE injects invalid SMILES and adds extraneous caption details while MolEdit outputs more valid SMILES and on-target captions.

|                   | M-Cap | F-Cap | M-Mol | F-Mol |
|-------------------|-------|-------|-------|-------|
| <b>Accuracy</b>   | 0.827 | 0.772 | 0.635 | 0.529 |
| $\Delta \uparrow$ | 65.5% | 54.4% | 90.7% | 58.9% |

Table 4: Accuracy of Expertise-Aware Editing Switcher under different settings.  $\Delta \uparrow$  represents the improvement in accuracy compared to switchers that do not consider expertise. For abbreviations, "M" denotes MoMu, "F" denotes MolFM; "Cap" denotes editing caption generation, "Mol" denotes editing molecule generation.

such as Structure set ("two hydroxy groups on the C-30 side-chain are located at positions 19 and 20.") and Origin set ("It derives from a D-mannitol."). By evaluating fine-tuning performance on this dataset (Figure 5), we observe that editing sensitivities vary across expertise domains, with structure edits being the most challenging. This divergence underscores the importance of adopting expertise-wise editing strategies. (2) **The effectiveness of expertise-wise editing.** We then examine whether MEKA actually learns to separate and route expertise during training by analyzing the expert activation distribution of the MoE, as shown in Figure 6. The results indicate that all experts consistently exhibit high activation rates, demonstrating that the MoE effectively isolates and routes different expertise during editing rather than collapsing onto a single route.

Next, we validate the effectiveness of expertise-aware editing switcher (EAES) in selectively activating relevant knowledge while suppressing untargeted information. As shown in Table 4, EAES achieves significantly higher switching accuracy than non-expertise-aware alternatives, with improvements of up to 90.7%. By validating these principles, we demonstrate that MolEdit’s modular design effectively addresses molecule-specific editing challenges, resulting in enhanced performance.

## 5.5 Case Study

We present case studies for two tasks using MolFM, comparing the outputs of GRACE and MolEdit on two representative examples in Fig 7. For molecule generation, GRACE frequently introduces extraneous or invalid tokens into the SMILES strings, such as incorrect brackets and the non-SMILES token "adopt" (highlighted in red). For caption generation, GRACE also hallucinates irrelevant caption details (e.g., "major species at pH 7.3" and "enol"), likely remnants from incomplete deletion of prior descriptions. In contrast, MolEdit produces outputs that closely align with the ground truth with minimal spillover, highlighting its robustness and precision in knowledge editing.

## 6 Conclusion

In this paper, we make the first attempt to edit MoLMs for both molecule and caption generation. To address the unique challenges associated with molecular editing, we propose MolEdit, a novel framework which consists of two key components: (1) MEKA, which directs molecular knowledge to specialized editing experts, enabling fine-grained control over multi-faceted updates; and (2) EAES, which maintains a memory bank of edited molecular expertise to activate MEKA only for highly relevant inputs. To enable systematic evaluation, we introduce MEBench, a comprehensive benchmark that assesses three critical dimensions: Reliability, Locality, and Generality. Extensive experiments demonstrate the effectiveness of MolEdit, prompting future work on editing MoLMs.

## Acknowledgment

This work is supported in part by the National Science Foundation (NSF) under grants IIS-2144209, IIS-2223769, CNS-2154962, BCS2228534, and CMMI-2411248; the Office of Naval Research (ONR) under grant N000142412636; and the Commonwealth Cyber Initiative (CCI) under grant VV-1Q24-011, and the gift funding from Netflix and Snap.

## Ethical Considerations

This work focuses on methodology for correcting errors in knowledge stored in molecule language models using public datasets and benchmarks. While we see limited ethical risks, potential concerns include misuse of editing to amplify unsafe molecular knowledge and inadvertent propagation of dataset biases. To mitigate these potential problems, we restrict evaluations to benign data, where the curated benchmark only provide benign targets. Overall, we view the ethical risks as minimal and manageable under standard responsible-use practices.

## References

- [1] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*. 65–72.
- [2] He Cao, Zijing Liu, Xingyu Lu, Yuan Yao, and Yu Li. 2023. Instructmol: Multimodal integration for building a versatile and reliable molecular assistant in drug discovery. *arXiv preprint arXiv:2311.16208* (2023).
- [3] Siyuan Cheng, Bozhong Tian, Qingbin Liu, Xi Chen, Yongheng Wang, Huajun Chen, and Ningyu Zhang. 2023. Can We Edit Multimodal Large Language Models?. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 13877–13888. doi:10.18653/v1/2023.emnlp-main.856
- [4] Siyuan Cheng, Bozhong Tian, Qingbin Liu, Xi Chen, Yongheng Wang, Huajun Chen, and Ningyu Zhang. 2023. Can we edit multimodal large language models? *arXiv preprint arXiv:2310.08475* (2023).
- [5] Siyuan Cheng, Ningyu Zhang, Bozhong Tian, Xi Chen, Qingbin Liu, and Huajun Chen. 2024. Editing language model-based knowledge graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17835–17843.
- [6] Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo Manica. 2023. Unifying molecular and textual representations via multi-task language modelling. In *International Conference on Machine Learning*. PMLR, 6140–6157.
- [7] Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164* (2021).
- [8] Jingcheng Deng, Zihao Wei, Liang Pang, Hanxing Ding, HuaWei Shen, and Xueqi Cheng. 2024. UnKE: Unstructured Knowledge Editing in Large Language Models. *arXiv preprint arXiv:2405.15349* (2024).
- [9] Yifan Deng, Spencer S Ericksen, and Anthony Gitter. 2024. Chemical Language Model Linker: blending text and molecules with modular adapters. *arXiv preprint arXiv:2410.20182* (2024).
- [10] Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui, and Lei Li. 2022. Calibrating factual knowledge in pretrained language models. *arXiv preprint arXiv:2210.03329* (2022).
- [12] Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. 2022. Translation between molecules and natural language. *arXiv preprint arXiv:2204.11817* (2022).
- [13] Carl Edwards, Qingyun Wang, Lawrence Zhao, and Heng Ji. 2024. L+ m-24: Building a dataset for language+ molecules@ acl 2024. *arXiv preprint arXiv:2403.00791* (2024).
- [14] Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Huajun Chen. 2023. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. *arXiv preprint arXiv:2306.08018* (2023).
- [15] Haisong Gong, Qiang Liu, Shu Wu, and Liang Wang. 2024. Text-guided molecule generation with diffusion language model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 109–117.
- [16] Tom Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. 2024. Aging with grace: Lifelong model editing with discrete key-value adapters. *Advances in Neural Information Processing Systems* 36 (2024).
- [17] Baixiang Huang, Canyu Chen, Xiong Xiao Xu, Ali Payani, and Kai Shu. 2024. Can Knowledge Editing Really Correct Hallucinations? *arXiv preprint arXiv:2410.16251* (2024).
- [18] Han Huang, Haitian Zhong, Tao Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2024. Vlkeb: A large vision-language model knowledge editing benchmark. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [19] Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. 2023. Transformer-patcher: One mistake worth one neuron. *arXiv preprint arXiv:2301.09785* (2023).
- [20] Saeedeh Javadi. 2024. *Knowledge Editing in Large Language Model*. Ph. D. Dissertation. Politecnico di Torino.
- [21] Ge Lei, Ronan Docherty, and Samuel J Cooper. 2024. Materials science in the era of large language models: a perspective. *Digital Discovery* (2024).
- [22] Jiatong Li, Yunqing Liu, Wenqi Fan, Xiao-Yong Wei, Hui Liu, Jiliang Tang, and Qing Li. 2024. Empowering molecule discovery for molecule-caption translation with large language models: A chatgpt perspective. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [23] Sihang Li, Zhiyuan Liu, Yanchen Luo, Xiang Wang, Xiangnan He, Kenji Kawaguchi, Tat-Seng Chua, and Qi Tian. 2024. Towards 3d molecule-text interpretation in language models. *arXiv preprint arXiv:2401.13923* (2024).
- [24] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. 2021. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208* (2021).
- [25] Zhoubo Li, Ningyu Zhang, Yunzhi Yao, Mengru Wang, Xi Chen, and Huajun Chen. 2023. Unveiling the pitfalls of knowledge editing for large language models. *arXiv preprint arXiv:2310.02129* (2023).
- [26] Chang Liao, Yemin Yu, Yu Mei, and Ying Wei. 2024. From Words to Molecules: A Survey of Large Language Models in Chemistry. *arXiv preprint arXiv:2402.01439* (2024).
- [27] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.
- [28] Pengfei Liu, Yiming Ren, Jun Tao, and Zhixiang Ren. 2024. Git-mol: A multimodal large language model for molecular science with graph, image, and text. *Computers in biology and medicine* 171 (2024), 108073.
- [29] Pengfei Liu, Jun Tao, and Zhixiang Ren. 2024. Scientific language modeling: A quantitative review of large language models in molecular science. *arXiv preprint arXiv:2402.04119* (2024).
- [30] Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Animashree Anandkumar. 2023. Multi-modal molecule structure-text model for text-based retrieval and editing. *Nature Machine Intelligence* 5, 12 (2023), 1447–1457.
- [31] Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, and Jian Tang. 2021. Pre-training molecular graph representation with 3d geometry. *arXiv preprint arXiv:2110.07728* (2021).
- [32] Xianggen Lv, Yan Guo, Haoran Li, Jin Liu, Shudong Huang, Bowen Ke, and Jiancheng Lv. 2024. DrugLLM: Open Large Language Model for Few-shot Molecule Generation. *arXiv preprint arXiv:2405.06690* (2024).
- [33] Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. 2023. MolCA: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. *arXiv preprint arXiv:2310.12798* (2023).
- [34] Yizhen Luo, Kai Yang, Massimo Hong, Xingyi Liu, and Zaiqing Nie. 2023. Molfm: A multimodal molecular foundation model. *arXiv preprint arXiv:2307.09484* (2023).
- [35] Yizhen Luo, Kai Yang, Massimo Hong, Xing Yi Liu, Zikun Nie, Hao Zhou, and Zaiqing Nie. 2024. Learning multi-view molecular representations with structured and unstructured knowledge. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2082–2093.
- [36] Yaohui Ma, Xiaopeng Hong, Shizhou Zhang, Huiyun Li, Zhilin Zhu, Wei Luo, and Zhiheng Ma. 2025. ComprehendEdit: A Comprehensive Dataset and Evaluation Framework for Multimodal Knowledge Editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 19323–19331.
- [37] Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. Memory-assisted prompt editing to improve GPT-3 after deployment. *arXiv preprint arXiv:2201.06009* (2022).
- [38] Vittorio Mazzia, Alessandro Pedrani, Andrea Caciolai, Kay Rottmann, and Davide Bernardi. 2024. A survey on knowledge editing of neural networks. *IEEE Transactions on Neural Networks and Learning Systems* (2024).
- [39] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems* 35 (2022), 17359–17372.
- [40] Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2022. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229* (2022).
- [41] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2021. Fast model editing at scale. *arXiv preprint arXiv:2110.11309* (2021).
- [42] Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. 2022. Memory-based model editing at scale. In *International Conference on Machine Learning*. PMLR, 15817–15831.
- [43] Shiwen Ni, Dingwei Chen, Chengming Li, Xiping Hu, Ruifeng Xu, and Min Yang. 2023. Forgetting before learning: Utilizing parametric arithmetic for knowledge updating in large language models. *arXiv preprint arXiv:2311.08011* (2023).
- [44] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the*

- 40th annual meeting of the Association for Computational Linguistics. 311–318.
- [45] Qizhi Pei, Lijun Wu, Kaiyuan Gao, Jinhua Zhu, Yue Wang, Zun Wang, Tao Qin, and Rui Yan. 2024. Leveraging Biomolecule and Natural Language through Multi-Modal Learning: A Survey. *arXiv preprint arXiv:2403.01528* (2024).
  - [46] Qizhi Pei, Lijun Wu, Kaiyuan Gao, Jinhua Zhu, and Rui Yan. 2024. 3D-MoIT5: Towards Unified 3D Molecule-Text Modeling with 3D Molecular Tokenization. *arXiv preprint arXiv:2406.05797* (2024).
  - [47] Qizhi Pei, Wei Zhang, Jinhua Zhu, Kehan Wu, Kaiyuan Gao, Lijun Wu, Yingce Xia, and Rui Yan. 2023. Biot5: Enriching cross-modal integration in biology with chemical knowledge and natural language associations. *arXiv preprint arXiv:2310.07276* (2023).
  - [48] Anton Sinitsin, Vsevolod Plokhhotnyuk, Dmitriy Pyrkin, Sergei Popov, and Artem Babenko. 2020. Editable neural networks. *arXiv preprint arXiv:2004.00345* (2020).
  - [49] Patrick Soga, Zhenyu Lei, Camille Bilodeau, and Jundong Li. 2025. Deep Interactions for Multimodal Molecular Property Prediction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 329–341.
  - [50] Bing Su, Dazhao Du, Zhao Yang, Yujie Zhou, Jiangmeng Li, Anyi Rao, Hao Sun, Zhiwu Lu, and Ji-Rong Wen. 2022. A molecular multimodal foundation model associating molecule graphs with natural language. *arXiv preprint arXiv:2209.05481* (2022).
  - [51] Raymond R Tice, Christopher P Austin, Robert J Kavlock, and John R Bucher. 2013. Improving the human hazard characterization of chemicals: a Tox21 update. *Environmental health perspectives* 121, 7 (2013), 756–765.
  - [52] Peng Wang, Zexi Li, Ningyu Zhang, Ziwen Xu, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2024. WISE: Rethinking the Knowledge Memory for Lifelong Model Editing of Large Language Models. *arXiv preprint arXiv:2405.14768* (2024).
  - [53] Renzhi Wang and Piji Li. 2024. Lemoe: Advanced mixture of experts adaptor for lifelong model editing of large language models. *arXiv preprint arXiv:2406.20030* (2024).
  - [54] Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. 2024. Knowledge editing for large language models: A survey. *Comput. Surveys* 57, 3 (2024), 1–37.
  - [55] John Wieting, Taylor Berg-Kirkpatrick, Kevin Gimpel, and Graham Neubig. 2019. Beyond BLEU: training neural machine translation with semantic similarity. *arXiv preprint arXiv:1909.06694* (2019).
  - [56] Yi Xiao, Xiangxin Zhou, Qiang Liu, and Liang Wang. 2024. Bridging Text and Molecule: A Survey on Multimodal Frameworks for Molecule. *arXiv preprint arXiv:2403.13830* (2024).
  - [57] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018).
  - [58] Li Yujian and Liu Bo. 2007. A normalized Levenshtein distance metric. *IEEE transactions on pattern analysis and machine intelligence* 29, 6 (2007), 1091–1095.
  - [59] Zhen Zeng, Leijiang Gu, Xun Yang, Zhangling Duan, Zenglin Shi, and Meng Wang. 2024. Visual-oriented fine-grained knowledge editing for multimodal large language models. *arXiv preprint arXiv:2411.12790* (2024).
  - [60] Zheni Zeng, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2022. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications* 13, 1 (2022), 862.
  - [61] Binchi Zhang, Zhengzhang Chen, Zaiyi Zheng, Jundong Li, and Haifeng Chen. 2025. Resolving editing-unlearning conflicts: A knowledge codebook framework for large language model updating. *arXiv preprint arXiv:2502.00158* (2025).
  - [62] Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi, Shengyu Mao, Jintian Zhang, Yuansheng Ni, et al. 2024. A comprehensive study of knowledge editing for large language models. *arXiv preprint arXiv:2401.01286* (2024).
  - [63] Qiang Zhang, Keyan Ding, Tianwen Lv, Xinda Wang, Qingyu Yin, Yiwen Zhang, Jing Yu, Yuhao Wang, Xiaotong Li, Zhuoyi Xiang, et al. 2024. Scientific large language models: A survey on biological & chemical domains. *Comput. Surveys* (2024).
  - [64] Yikun Zhang, Geyan Ye, Chaohao Yuan, Bo Han, Long-Kai Huang, Jianhua Yao, Wei Liu, and Yu Rong. 2024. Atomas: Hierarchical Alignment on Molecule-Text for Unified Molecule Understanding and Generation. *arXiv preprint arXiv:2404.16880* (2024).
  - [65] Haiteng Zhao, Shengchao Liu, Ma Chang, Hannan Xu, Jie Fu, Zhihong Deng, Lingpeng Kong, and Qi Liu. 2023. Gimlet: A unified graph-text model for instruction-based molecule zero-shot learning. *Advances in Neural Information Processing Systems* 36 (2023), 5850–5887.
  - [66] Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. 2023. Can we edit factual knowledge by in-context learning? *arXiv preprint arXiv:2305.12740* (2023).
  - [67] Zexuan Zhong, Zhengxuan Wu, Christopher D Manning, Christopher Potts, and Danqi Chen. 2023. Mquake: Assessing knowledge editing in language models via multi-hop questions. *arXiv preprint arXiv:2305.14795* (2023).