

Enhancing Chemical Reaction and Retrosynthesis Prediction with Large Language Model and Dual-task Learning

Xuan Lin¹, Qingrui Liu¹, Hongxin Xiang^{2*}, Daojian Zeng³ and Xiangxiang Zeng²

¹School of Computer Science, Xiangtan University

²College of Computer Science and Electronic Engineering, Hunan University

³Institute of AI and Targeted International Communication, Hunan Normal University

Abstract

Chemical reaction and retrosynthesis prediction are fundamental tasks in drug discovery. Recently, large language models (LLMs) have shown potential in many domains. However, directly applying LLMs to these tasks faces two major challenges: (i) lacking a large-scale chemical synthesis-related instruction dataset; (ii) ignoring the close correlation between reaction and retrosynthesis prediction for the existing fine-tuning strategies. To address these challenges, we propose ChemDual, a novel LLM framework for accurate chemical synthesis. Specifically, considering the high cost of data acquisition for reaction and retrosynthesis, ChemDual regards the reaction-and-retrosynthesis of molecules as a related recombination-and-fragmentation process and constructs a large-scale of 4.4 million instruction dataset. Furthermore, ChemDual introduces an enhanced LLaMA, equipped with a multi-scale tokenizer and dual-task learning strategy, to jointly optimize the process of recombination and fragmentation as well as the tasks between reaction and retrosynthesis prediction. Extensive experiments on Mol-Instruction and USPTO-50K datasets demonstrate that ChemDual achieves state-of-the-art performance in both predictions of reaction and retrosynthesis, outperforming the existing conventional single-task approaches and the general open-source LLMs. Through molecular docking analysis, ChemDual generates compounds with diverse and strong protein binding affinity, further highlighting its strong potential in drug design. Code and appendix are available at <https://github.com/JacklinGroup/ChemDual>.

1 Introduction

Chemical reaction and retrosynthesis prediction are fundamental tasks in many fields of chemistry, forming the backbone of synthetic route design, compound optimization, and drug discovery [Ucak *et al.*, 2022]. These tasks involve

predicting the outcomes of chemical reactions and identifying potential synthetic pathways to create target molecules. Traditional approaches have relied heavily on expert knowledge and manual analysis [Liang *et al.*, 2024], which are time-consuming and resource-limiting. With the advancement of computational technologies, automated methods have emerged as powerful alternatives to offer significant improvements in speed and efficiency for these critical tasks.

Large language models (LLMs) have gained more attention in various domains [OpenAI, 2024], they leverages advanced natural language processing techniques to process and analyze complex biochemical data [Zhao *et al.*, 2025]. Recently, Mol-Instruction [Fang *et al.*, 2024] has show promising results in addressing the inherent challenges of chemical reaction and retrosynthesis prediction. Despite these advancements, LLMs still face the following two major challenges, which limit their accuracy in predictions of reaction and retrosynthesis compared with traditional machine learning models [Guo *et al.*, 2023].

Challenge 1: Lacking a large-scale chemical synthesis-related instruction dataset. The success of LLM depends on the availability of large and high-quality data [Dai *et al.*, 2019]. In predictions of reaction and retrosynthesis, high-quality data comes from actual chemical synthesis experiments in the laboratory, so it is costly and inefficient, resulting in limited data size. Therefore, this motivates us to find an alternative with low acquisition cost (capable of quickly generating large-scale data) and high data quality (high relevance to real chemical synthesis). As shown in the top of Figure 1(a), we find that fragments generated by breaking of retrosynthetically interesting chemical substructures (BRICS) [Degen *et al.*, 2008] are highly correlated with reactants. We also show the similarity distribution of all molecules in the bottom of Figure 1(a), which has a global average similarity of 66.5%, indicating that we can use BRICS to generate a large amount of data to assist in the learning of predictions of reaction and retrosynthesis. Therefore, we construct a large-scale chemical synthesis-related instruction dataset, which includes 4.4 million molecules and the corresponding fragments extracted by BRICS.

Challenge 2: Ignoring the correlation between predictions of reaction and retrosynthesis. As shown in the left sub-figure of Figure 1(b), existing methods treat Molecule-to-Reactants and Reactants-to-Molecules as two indepen-

* Corresponding author (xianghx@hnu.edu.cn)

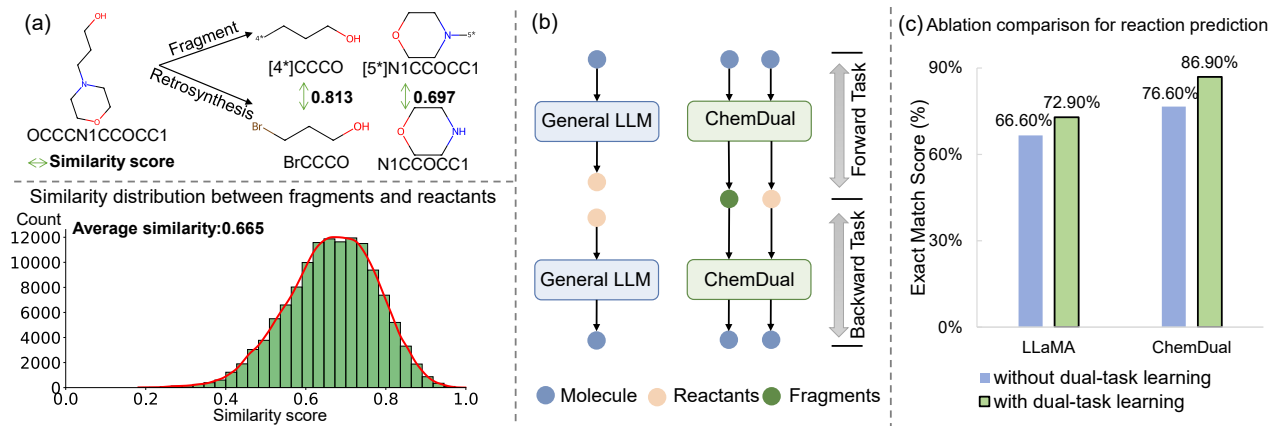


Figure 1: (a) The upper subgraph is an example of the similarity between fragments and reactants for a molecule "OCCCN1CCOCC1". The lower subgraph is the similarity distribution between fragments and reactants for all molecules. (b) Single-task learning paradigm (left sub-figure) and dual-task learning paradigm of ChemDual (right sub-figure). (c) The Exact Match score (%) on the chemical reaction prediction task by using or not using the dual-task learning.

dent tasks for learning [Wang *et al.*, 2023; Liang *et al.*, 2023]. However, Molecule-to-Reactants and Reactants-to-Molecules are inverse processes of each other, ignoring the correlation between them will limit LLM in understanding chemical synthesis processes [He *et al.*, 2016]. Therefore, we regard Molecule-to-Reactants and Reactants-to-Molecules as forward task and backward task respectively and use a dual-task learning strategy to optimize these two tasks simultaneously. As shown in Figure 1(c), the LLaMA equipped with dual-task learning achieves 6.3% performance improvement of Exact Match score on reaction prediction task.

Inspired by these, we propose ChemDual, an enhanced LLaMA equipped with a multi-scale tokenizer and dual-task learning strategy. The multi-scale tokenizer is used to enhance the ability of the model in capturing molecular structures at different scales (such as dummy atom, functional group and fragment, etc.). In dual-task learning, as shown in the right sub-figure of Figure 1(b), we pretrain on molecule-to-fragments and fragments-to-molecule dual tasks constructed from 4.4 million data to enhance ChemDual in understanding of general chemical synthesis, and further fine-tune ChemDual to perform task-specific predictions using molecule-to-reactants and reactants-to-molecule dual tasks. In Figure 1(c), we find that ChemDual achieves better performance compared to LLaMA. The main contributions of this work are as follows:

- We construct a large-scale database, which includes 4.4 million molecules and the corresponding fragments, for learning general chemical synthesis-related knowledge.
- We propose an enhanced LLM framework based on LLaMA, called ChemDual, equipped with a multi-scale tokenizer and a dual-task learning strategy for learning informative molecular representation from different scales of structures and capturing the correlation between forward and backward processes, respectively.
- Extensive experiments on Mol-Instruction and USPTO-50K show that ChemDual outperforms existing models

in predictions of reaction and retrosynthesis. In addition, Case studies demonstrate that ChemDual can generate compounds with favorable binding properties and specific molecular interactions.

2 Related Work

Recent researches have proposed a large number of large language models (LLMs), such as LLaMA [Touvron *et al.*, 2023] and ChatGLM [GLM *et al.*, 2024]. Given the important role of chemical reaction and retrosynthesis prediction tasks in drug discovery, many chemical-based domain-specific models have been further proposed [Taylor *et al.*, 2022; Pei *et al.*, 2023]. The core idea of these models is to use reaction- or retrosynthesis-based instruction datasets to further fine-tune the existing LLMs to enable them to have predictive capabilities for reaction or retrosynthesis. BioT5+ [Pei *et al.*, 2024] involves molecule a dataset of 30k SELFIES [Krenn *et al.*, 2020] mapped to IUPAC names, with multi-task instruction-based fine-tuning. Text+Chem T5 [Christofidellis *et al.*, 2023] and Mol-Instruction [Fang *et al.*, 2024] specifically address chemical tasks by integrating cross-domain knowledge into their frameworks, utilizing a dataset of about 30k samples. Retroformer [Wan *et al.*, 2022], trained on the augmentative USPTO-50k [Schneider *et al.*, 2016] dataset containing more than 50k reaction samples with a single-task learning strategy, has inspired other works [Schwaller *et al.*, 2019] [Seo *et al.*, 2021] that improve performance by incorporating a variety of reasonable SMILES data enhancement methods for the same molecule under a single-task learning framework. Different from previous methods, we cheaply and efficiently construct an instruction fine-tuning database with 4.4 million molecules related to chemical reaction and retrosynthesis for learning general synthesis-related knowledge, and introduce a dual-task learning strategy to simultaneously capture the forward and backward relationships when learning synthesis.

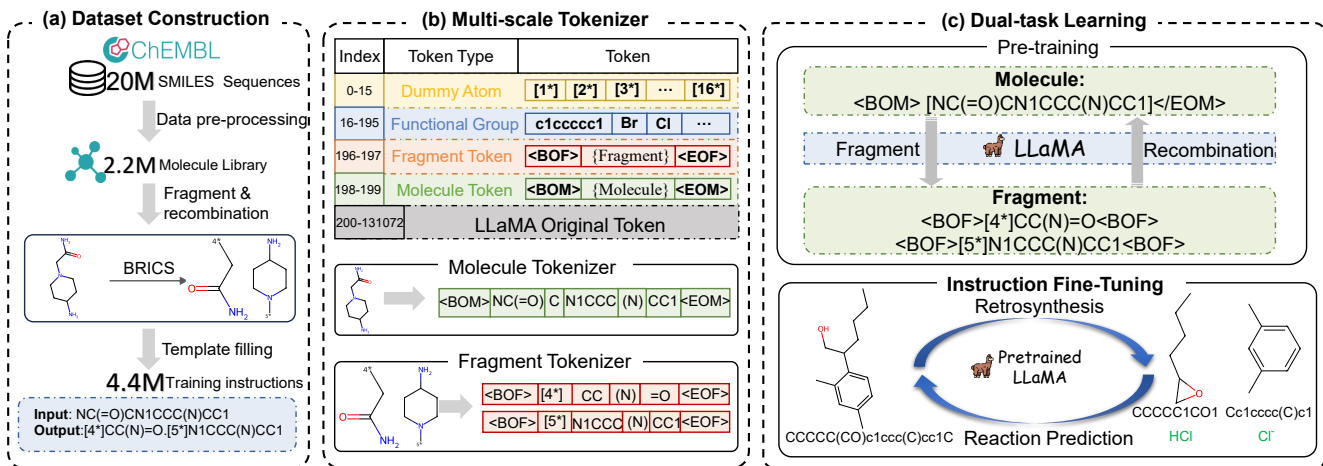


Figure 2: The overall of proposed ChemDual.

3 Method

The overview of our proposed ChemDual is shown in Figure 2, which is divided into three main modules. In the dataset construction module, we construct a 4.4M Molecule-Fragments database from 20M SMILES sequences with breaking of retrosynthetically interesting chemical substructures (BRICS) [Degen *et al.*, 2008] (Section 3.1). Subsequently, in the module of multi-scale tokenizer, we extend the tokenizer the existing tokenizer to capture molecular information at different scales such as reaction and fragment (Section 3.2). Finally, in the module of dual-task learning, we use dual-task learning on molecules and fragment as well as reaction and retrosynthesis to help LLM learn informative representation (Section 3.3).

3.1 Dataset Construction

As shown in Figure 2(a), we construct a chemical synthesis-related instruction dataset through data pre-processing, fragment-recombination operation, and template filling.

Data pre-processing. We collected 20M molecular SMILES sequences from the ChEMBL-34 database [Zdrzil *et al.*, 2024] and preprocessed these molecules according to three criteria: (i) removing duplicates; (ii) filtering out invalid molecules using RDKit [Bento *et al.*, 2020]; (iii) excluding molecules with weight greater than 1000. Next, we tokenized the molecular SMILES sequences and removed sequences with a token length of more than 512 to prevent exceeding the maximum token limit of LLM. Finally, we obtained a molecule library containing 2.2M SMILES sequences.

Fragment-recombination operation. Based on the molecular library, we adopt the BRICS algorithm to generate multiple fragments for the SMILES of the molecule and use these fragments to recombine back to the original SMILES. We regard the process of molecular fragmentation and recombination as a dual task. Meanwhile, we implement an adaptive fragmentation approach to alleviate memory overflow caused by excessive fragmentation when long SMILES sequences is processed by BRICS. Specifically, we adjust the number of

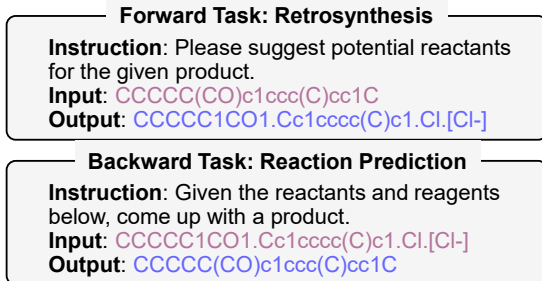


Figure 3: Examples of instruction of fine-tuning dataset.

fragments based on the length of molecule SMILES, and the maximum number of fragments is determined as follows:

$$n = \begin{cases} L, & \text{if } L < k \\ \min \left(L, \left\lceil \frac{L}{k} \right\rceil^\alpha \right), & \text{otherwise} \end{cases} \quad (1)$$

where L represents the length of the SMILES sequence, α denotes the elasticity factor to accommodate the capacities of different machine memory, and k is the average length of SMILES sequences in molecule library. Interestingly, we observe that these fragments can be reassembled into a molecule, we term it as *recombination*. In BRICS algorithm, dummy atoms are introduced at cleavage sites during the fragment labeling. We do not apply any additional processing to these dummy atoms during fragmentation, while these dummy atoms are removed and replaced with carbon (C) atoms to prevent the generation of invalid SMILES strings during recombination.

Template filling. To obtain more task-specific data and to adapt to the strong instruction-following abilities of LLMs, we design templates for chemical reaction and retrosynthesis prediction tasks. As shown in Figure 3, in the forward task (i.e., retrosynthesis), we fix a general question format and then extract the molecular SMILES from molecule library to fill the description part of a predefined template, result-

ing in a natural question as instructions. The backward task (i.e., reaction prediction) is to take multiple products and corresponding reaction types (marked in green) as the question template. Through the entire pipeline, we finally construct a dataset containing 4.4M fine-tuning instructions for downstream prediction tasks.

3.2 Multi-scale Tokenizer

BioT5 [Pei *et al.*, 2023] has demonstrated the advantages of using specialized tokenizer for LLMs to improve task-specific performance. As shown in Figure 2(b), our proposed ChemDual further extends the tokenizer of LLaMA 3.1 by incorporating three additional types of tokenization strategies to handle the complexities of chemical data. The first category includes dummy atoms ([1*], [2*], [3*], ..., [16*]), which represent molecular fragments generated through the BRICS algorithm and encompass a total of 16 distinct dummy atom types. The second category consists of 180 functional groups commonly found in chemical structures, such as benzene rings, halogen atoms, and other commonly encountered functional groups in chemical structures. The third category introduces 4 special tokens, <BOF> and <EOF> for molecule and <BOM> and <EOM> for fragment, allowing the model to distinguish whether a sequence is a complete molecule or the corresponding fragment. For a set of molecules or fragments, the input and output sequences are typically separated by a dot ("."). However, during tokenization, these sequences are further split into individual components using the special tokens discussed earlier.

3.3 Dual-task Learning

As illustrated in Figure 2(c), we introduce a dual-task learning in both pre-training and instruction fine-tuning, to address the insufficient understanding of LLMs to the inherent structural information of SMILES. By incorporating dual information into model training process, we aim to enable LLMs to capture more implicit chemical knowledge, thereby enhancing their understanding of chemical structure. Specifically, our dual tasks include forward tasks and backward tasks, which use (fragmentation, recombination) and (reaction, retrosynthesis) to form the forward task and the backward task pair. This design is analogous to English–French translation and back-translation in NLP, where bidirectional learning reinforces semantic understanding. The dual tasks can be represented mathematically as follows:

$$f_{\text{forward}} : X \rightarrow Y, \quad (X, Y) \in \{(M, \mathcal{F}), (P, \mathcal{R})\}, \quad (2)$$

$$f_{\text{backward}} : Y \rightarrow X, \quad (Y, X) \in \{(\mathcal{F}, M), (\mathcal{R}, P)\}, \quad (3)$$

where M , P , \mathcal{F} and \mathcal{R} represent a molecule, the reaction product, the molecular fragments and reactants, respectively.

The dual scenario can be represented by a more generalized function, we consider the joint probability of a data pair (x, y) , where $x \in X$ and $y \in Y$. Let $P(x)$ and $P(y)$ indicate the marginal distribution of x , y respectively, $P(y|x; LLM)$ indicate the conditional probability of generating y from x using the LLM, and $P(x|y; LLM)$ indicate the conditional probability of generating x from y via LLM. And we have:

$$P(x, y) = P(x)P(y|x; LLM) = P(y)P(x|y; LLM). \quad (4)$$

With the proposed dual-task learning framework, we can define the training mechanism, loss function, and optimization objective for LLM. Let $\mathcal{D} = (x_i, y_i)_{i=1}^N$ represent the dataset used for training, where x_i represents the input sequence, and y_i denotes the target sequence. Given an input sequence, the model is used to predict the probability distribution of possible output tokens, and its optimization goal is to maximize the likelihood of the correct sequence y_i . The probability distribution of the output vocabulary \mathcal{V} of the model at each time step is formalized as follows:

$$P(y_i | x_i; \theta) = \text{softmax}(z_i), \quad (5)$$

where z_i denotes the logits predicted by the model for the input sequence x_i , and θ represents the model parameters. We adopt cross-entropy loss as the objective function to measure the divergence between the predicted probability distribution and the true distribution (one-hot encoded label). The cross-entropy loss for a single training example (x_i, y_i) is defined as follows:

$$\mathcal{L}(x_i, y_i; \theta) = - \sum_{j=1}^{|\mathcal{V}|} \mathbb{I}(y_i = j) \log P(y_i = j | x_i; \theta), \quad (6)$$

where $\mathbb{I}(y_i = j)$ is the indicator function that equals 1 when y_i is the true label, and 0 otherwise.

Pre-training. We perform pre-training on the fragment-recombination tasks. Specifically, we define fragment as the forward task and recombination as the backward task. The pre-training process of ChemDual is optimized as follows:

$$\mathcal{L}_{pt}(x, y; \theta) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(x_i, y_i; \theta), \quad (7)$$

where $(x_i, y_i) \in \{(M, \mathcal{F}), (\mathcal{F}, M)\}$ evaluates the performance of fragment and recombination prediction on pre-training dataset. Dual-task learning enables the model to capture intrinsic relationships between molecular components, crucial for reaction and retrosynthesis prediction.

Instruction fine-tuning. We adapt the dual-task learning to fine-tuning more specific chemical synthesis tasks. Specifically, we regard retrosynthesis as the forward task, and reaction prediction is serves as the backward task. The fine-tuning process of ChemDual is optimized as follows:

$$\mathcal{L}_{ft}(x, y; \theta) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(x_i, y_i; \theta), \quad (8)$$

where $(x_i, y_i) \in \{(M, \mathcal{R}), (\mathcal{R}, M)\}$ accesses the performance of downstream tasks on the training dataset.

4 Experiments and Results

In this section, we first introduce compare our model ChemDual with baselines on reaction and retrosynthesis prediction. After that, we investigate the effectiveness of instruction dataset and fine-tuning strategy in our model. Finally, we conduct the visualization analysis and case study via molecular docking. Details of the experimental settings are provided in Appendix A.

Model	EXACT \uparrow	BLEU \uparrow	LEVENSHTEIN \downarrow	RDKit FTS \uparrow	MACCS FTS \uparrow	MORGAN FTS \uparrow	VALIDITY \uparrow
Reaction Prediction							
Alpaca*	0.000	0.065	41.989	0.004	0.024	0.008	0.138
Baize*	0.000	0.044	41.500	0.004	0.025	0.009	0.097
ChatGLM*	0.000	0.183	40.008	0.050	0.100	0.044	0.108
LLaMA*	0.000	0.020	42.002	0.001	0.002	0.001	0.039
Vicuna*	0.000	0.057	41.690	0.007	0.016	0.006	0.059
Galactica*	0.000	0.468	35.021	0.156	0.257	0.097	0.946
Text+Chem T5*	0.239	0.782	20.413	0.705	0.789	0.652	0.762
Mol-Instruction*	0.503	0.883	13.41	0.756	0.863	0.708	1.000
BioT5+	<u>0.864</u>	0.993	<u>3.403</u>	<u>0.949</u>	<u>0.975</u>	<u>0.935</u>	1.000
ChemDual	0.869	<u>0.991</u>	2.099	0.964	0.980	0.956	<u>0.996</u>
Retrosynthesis							
Alpaca*	0.000	0.063	46.915	0.005	0.023	0.007	0.160
Baize*	0.000	0.095	44.714	0.025	0.050	0.023	0.112
ChatGLM*	0.000	0.117	48.365	0.056	0.075	0.043	0.046
LLaMA*	0.000	0.036	46.844	0.018	0.029	0.017	0.010
Vicuna*	0.000	0.057	46.877	0.025	0.030	0.021	0.017
Galactica*	0.000	0.452	34.940	0.167	0.274	0.134	0.986
Text+Chem T5*	0.141	0.765	24.043	0.685	0.765	0.585	0.698
Mol-Instruction*	0.333	0.842	17.642	0.704	0.815	0.646	1.000
BioT5+	0.642	0.969	6.710	0.897	0.930	0.866	1.000
ChemDual	0.670	0.976	6.516	0.901	0.933	0.893	<u>0.995</u>

Table 1: Performance comparison for reaction and retrosynthesis prediction on the Mol-Instruction dataset. [*] indicates the results are taken from Mol-Instruction [Fang *et al.*, 2024]. The **best** and suboptimal results are shown in bold and underline.

Model	Top-1 (%)	Top-3 (%)	Top-5 (%)	Top-10 (%)
Transformer	42.40	58.60	63.80	67.70
BioT5+	<u>44.40</u>	<u>59.56</u>	<u>61.32</u>	<u>73.43</u>
InstructMol	30.15	51.72	57.12	64.91
ChemDual	46.25	60.14	66.95	77.42
Retroformer	47.89	62.87	66.59	70.68
ChemDual*	49.95	67.67	70.52	78.31
Improvement(%)	2.06	4.79	3.93	7.63

Table 2: Performance comparison on USPTO-50K dataset, [*] means ChemDual adopts the Retroformer’s transformer module.

4.1 Experiment Setup

Dataset. We use Mol-Instruction and ChemLLMBench dataset to evaluate ChemDual on reaction and retrosynthesis prediction. In addition, we evaluate the retrosynthesis task on the USPTO-50 dataset. The Mol-Instruction consists of 282,304 molecules paired instructions describing their reactions, and each sample is represented by a tuple (input, output), where the input could be either a set of reactants or a molecule, and the output corresponds to the predicted products or reactants. ChemLLMBench consists of a series of chemistry-related tasks. In this work, we evaluate ChemDual on the reaction and retrosynthesis prediction tasks within ChemLLMBench. We follow the same test set as used in ChemDFM [Zhao *et al.*, 2025]. The USPTO-50k dataset contains 50,016 reaction examples that are grouped into 10 reaction classes, and each sample represents a synthetic reaction step with corresponding reactants and products. We keep the same data split as Graph2SMILES [Tu and Coley, 2022].

Baseline models. We compare ChemDual against a variety of baselines which can be categorized as follows:

- **General LLMs:** Alpaca [Taori *et al.*, 2023], Baize [Xu *et al.*, 2023], ChatGLM [GLM *et al.*, 2024], LLaMA [Touvron *et al.*, 2023] and Vicuna [Chiang *et al.*, 2023],
- **Specialized models:** Galactica [Taylor *et al.*, 2022], Text+Chem T5 [Christofidellis *et al.*, 2023], Mol-Instruction, BioT5+ [Pei *et al.*, 2024], InstructMol [Cao *et al.*, 2025], Retroformer [Wan *et al.*, 2022].

Evaluation metrics. For reaction and retrosynthesis tasks, we employ various metrics to evaluate the effectiveness of ChemDual and all baselines, including EXACT, BLEU [Papineni *et al.*, 2002], LEVENSHTEIN [Levenshtein, 1966], three molecular fingerprints (FTS) similarity scores including RDKit [Schneider *et al.*, 2015], MACCS [Durant *et al.*, 2002], MORGAN [Rogers and Hahn, 2010] and VALIDITY score refers to whether the SMILES can be successfully processed by RDKit. Additionally, we further adopt Top-k accuracy as evaluation metric on retrosynthesis prediction.

4.2 Comparison Results

Table 1 shows the comparison results for reaction and retrosynthesis prediction on Mol-Instruction. Detailed results on ChemLLMBench can be found in Appendix B.

Results on reaction prediction. As shown in the top of Table 1, we observe that ChemDual achieves a substantial improvement with the highest EXACT score of 0.869 and the lowest LEVENSHTEIN distance of 2.099. Moreover, ChemDual consistently outperformed all baselines across multiple structure similarity metrics such as RDKit, MACCS, and




Figure 4: Comparison of ablation experiments for reaction prediction and retrosynthesis prediction on Mol-Instruction.

MORGAN FTS, with a promising scores of 0.964, 0.980, and 0.956, respectively. This indicates that our proposed ChemDual can perform highly chemical relevant predictions compared to other methods. However, ChemDual achieves competitive VALIDITY score with Mol-Instruction and BioT5+, the slight gap may be attributed to the reason that ChemDual adopt SMILES representations as input while Mol-Instruction and BioT5+ employ the unique SELFIES format.

Results on retrosynthesis prediction. The bottom of Table 1 shows that ChemDual achieves the best performance of 0.670 in EXACT, 0.976 in BLEU, and 6.516 in LEVENSHTEIN distance, respectively, with comparison to other models, including the strong baseline BioT5+. Furthermore, ChemDual also surpasses BioT5+ by 0.4%, 0.3%, and 2.7% in terms of RDK, MACCS, and MORGAN FTS, respectively, which implies that ChemDual performs the ability to generate accurate and chemical retrosynthesis predictions.

We further evaluate the retrosynthesis performance of ChemDual in USPTO-50K dataset. As shown in Table 2, ChemDual achieves the best performance in all top-k accuracy, which demonstrates its robustness and scalability across datasets. Specifically, ChemDual outperforms BioT5+ and other baseline models at least by 1.85%, 0.58%, 3.15%, and 3.99% in terms of Top-k ($k=1, 3, 5$, and 10), respectively. Compared with Retroformer, the fine-tuned ChemDual* consistently makes improvements of 2.06% in Top-1, 4.79% in Top-3, 3.93% in Top-5, and 7.63% in Top-10, respectively.

Results on fragment and recombination. To investigate whether ChemDual accomplish the molecular fragment and recombination tasks, we evaluate our proposed model using 1,000 samples from our constructed instruction test set. The details of results refer to Appendix C.

4.3 Ablation Study

To investigate the effectiveness of each component to our model, we consider the following variants of ChemDual. Figure 4 shows the results of ChemDual and its variants on reaction and retrosynthesis prediction.

- *LLaMA with dual-task learning* (w dual) removes the datasets related to molecule fragments and recombination, and we trained the LLaMA 3.1 only using Mol-Instruction dataset.
- *ChemDual/LLaMA without dual-task learning* (w/o dual) deletes the datasets related to reaction prediction in retrosynthesis prediction task, and the datasets related




Figure 5: Visualizations of LLaMa and ChemDual using the t-SNE.

to retrosynthesis is prediction are deleted reaction prediction task, respectively.

- *ChemDual without pre-training* (w/o pre) directly fine-tunes the LLaMA 3.1 on Mol-Instruction and our constructed instruction datasets with a ratio of 1:1.

Dataset ablation. ChemDual outperforms the LLaMA (w dual) on all metrics, which demonstrates the importance of leveraging our constructed instruction dataset for enhancing predictive performance.

Pre-training ablation. The comparison of ChemDual and ChemDual (w/o pre) reveals a consistent improvement on model performance for the pre-trained model (i.e., ChemDual) across reaction and retrosynthesis prediction tasks. This highlights the critical role of pre-training in capturing the intrinsic relationships between molecular component.

Dual-task learning ablation. Dual-task method (i.e., LLaMA (w dual)) performs better than single-task model (i.e., ChemDual/LLaMA (w/o dual)), which illustrates that the model is fine-tuned by the dual-task paradigm, either independently or with the inclusion of molecular fragment and recombination datasets, consistently outperformed than that is trained on a single dataset. These results underscore the effectiveness of the dual-task learning approach, which promotes the synergistic learning across tasks and improves the capabilities of model on chemical reasoning.

The ablation studies further validate the critical role of instruction dataset and dual-task learning paradigm in enhancing the performance of general LLM on reaction and retrosynthesis prediction tasks. We also highlight that a well-designed pre-training strategy allows LLM to better grasp the semantics of chemical structures, enabling them to make more precise predictions across multiple downstream tasks.




Figure 6: Visualization of the molecular docking complex between the top1 molecule generated by ChemDual and the target protein MAP2.

5 Analysis

Visualization. To understand the superiority of ChemDual over the general LLM (i.e., LLaMa), we employ t-SNE [van der Maaten and Hinton, 2008] to visualize the learned representations of input prompt. Since there are dozens of reaction types, we select 8 reaction types and each with 50 sample data for visualization. Figure 5 shows the visualization on USPTO-50K dataset. We clearly observe that input prompts are more tightly clustered in ChemDual compared with LLaMa, which implies that ChemDual can learn more contextual information by effectively capturing the chemical synthesis-related knowledge from our newly constructed instruction dataset. It is worth noting that more compact clusters in ChemDual on 8 reaction types can well illustrate the excellent ability of ChemDual on retrosynthesis prediction.

Molecule recombination and docking. We conduct a case study to investigate the ability of our proposed ChemDual on molecular recombination. Traditional methods struggle with balancing chemical validity and structural diversity. Here we further explore how our constructed instruction dataset addresses this challenge by generating molecules with high chemical similarity. Specifically, we select the compound (CHEMBL419018) from the available ChEMBL database. By fragmenting the compound with BRICS, we obtain multiple fragments as shown in Figure 6. Based on these fragments, we are able to generate novel compounds (the details refers to Figure 8 in Appendix D) via ChemDual. These generated compounds inherit characteristics from their original fragments and are highly similar to the original molecules in molecular fingerprints, which demonstrates the ability of ChemDual to maintain key chemical features while generating valid molecular structures, ensuring both chemical integrity and similarity during the recombination process.

Methionine aminopeptidase 2 (MAP2) has been recognized as a therapeutic target with several inhibitors, demonstrating the promising anti-cancer and anti-obesity activi-

ties in both preclinical and clinical studies. Therefore, we use the MAP2 as target protein (CHEMBL419018 in the ChEMBL34 dataset) and conduct docking experiments between the protein and the generated compounds to evaluate binding affinities. As shown in Figure 8, the binding affinity values of the generated compounds range from -6.3 to -8.4 kcal/mol, which indicates that the molecules generated by ChemDual have good binding interactions with the target protein. To gain deeper insights, we use AutoDock [Morris *et al.*, 2009] to further visualize the docking result and binding interactions of the first generated compound in Figure 6. The visualization highlights three key hydrogen bonds formed by the ligand: the first is between the O3 atom of the ligand and the 1HH2 atom of the ARG-417 residue; the second is between the O3 atom of the ligand and the HH atom of the TYR-383 residue; and the third is between the O1 atom of the ligand and the NH group of the LEU18 residue. These interactions underscore the stability of the ligand within the protein’s active site, demonstrating that ChemDual can not only recombine structurally similar molecules but also yield compounds with stable, specific binding characteristics in target interactions.

6 Conclusion

We propose ChemDual, an enhanced LLaMA-based LLM for chemical reaction and retrosynthesis prediction. It constructs a 4.4 million molecule database to alleviate the lack of a large-scale chemical synthesis-related instruction dataset, employs a multi-scale tokenizer to capture structural information, and introduces a dual-task learning strategy to jointly optimize the process of recombination and fragmentation as well as the tasks between reaction and retrosynthesis prediction. Experiments on Mol-Instruction and USPTO-50K demonstrate that ChemDual outperforms existing methods, while molecular docking results show its ability to generate compounds with strong affinity for target proteins, highlighting its potential in molecular design.

Acknowledgements

We would like to thank all anonymous reviewers for their valuable feedback and suggestions. This work was supported in part by the National Science and Technology Major Project (Grant No. 2023ZD0120902), in part by the National Natural Science Foundation of China (Grant No. 62202413, 62425204, 62122025, and 62276095), and in part by the Research Foundation of the Education Bureau of Hunan Province, China (Grant No. 23B0129).

References

- [Bento *et al.*, 2020] A Patrícia Bento, Anne Hersey, Eloy Félix, Greg Landrum, Anna Gaulton, Francis Atkinson, Louisa J Bellis, Marleen De Veij, and Andrew R Leach. An open source chemical structure curation pipeline using rdkit. *Journal of Cheminformatics*, 12:1–16, 2020.
- [Cao *et al.*, 2025] He Cao, Zijing Liu, Xingyu Lu, Yuan Yao, and Yu Li. InstructMol: Multi-modal integration for building a versatile and reliable molecular assistant in drug discovery. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 354–379, 2025.
- [Chiang *et al.*, 2023] Wei-Lin Chiang, Zhuohan Li, Zi Lin, et al. Vicuna: An open-source chatbot impressing GPT-4 with 90%* ChatGPT quality, 2023.
- [Christofidellis *et al.*, 2023] Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo Manica. Unifying molecular and textual representations via multi-task language modelling. In *Proceedings of the 41th International Conference on Machine Learning*, pages 6140–6157, 2023.
- [Dai *et al.*, 2019] Hanjun Dai, Chengtao Li, Connor W. Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. In *Proceedings of the 32th Annual Conference on Neural Information Processing Systems*, pages 8870–8880, 2019.
- [Degen *et al.*, 2008] Jorg Degen, Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. On the art of compiling and using ‘drug-like’ chemical fragment spaces. *ChemMedChem*, 3(10):1503, 2008.
- [Durant *et al.*, 2002] Joseph L Durant, Burton A Leland, Douglas R Henry, and James G Nourse. Reoptimization of mdl keys for use in drug discovery. *Journal of Chemical Information and Computer Sciences*, 42(6):1273–1280, 2002.
- [Fang *et al.*, 2024] Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Huajun Chen. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. In *Proceedings of the 12th International Conference on Learning Representations*, 2024.
- [GLM *et al.*, 2024] Team GLM, Aohan Zeng, and et al. ChatGLM: A family of large language models from GLM-130B to GLM-4 all tools, 2024.
- [Guo *et al.*, 2023] Taicheng Guo, Kehan Guo, Bozhao Nan, Zhenwen Liang, Zhichun Guo, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. What can large language models do in chemistry? A comprehensive benchmark on eight tasks. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 59662–59688, 2023.
- [He *et al.*, 2016] Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. Dual learning for machine translation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 820–828, 2016.
- [Krenn *et al.*, 2020] Mario Krenn, Florian Häse, Akshat Kumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, 2020.
- [Levenshtein, 1966] Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966.
- [Liang *et al.*, 2023] Ke Liang, Yue Liu, Sihang Zhou, Wenxuan Tu, Yi Wen, Xihong Yang, Xiangjun Dong, and Xinwang Liu. Knowledge graph contrastive learning based on relation-symmetrical structure. *IEEE Transactions on Knowledge and Data Engineering*, 36(1):226–238, 2023.
- [Liang *et al.*, 2024] Ke Liang, Lingyuan Meng, Yue Liu, Meng Liu, Wei Wei, Suyuan Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, and Xinwang Liu. Simple yet effective: Structure guided pre-trained transformer for multi-modal knowledge graph reasoning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1554–1563, 2024.
- [Morris *et al.*, 2009] Garrett M. Morris, Ruth Huey, William Lindstrom, Michel F. Sanner, Richard K. Belew, David S. Goodsell, and Arthur J. Olson. Autodock4 and autodock-tools4: Automated docking with selective receptor flexibility. *Journal of Computational Chemistry*, 30(16):2785–2791, 2009.
- [OpenAI, 2024] OpenAI. ChatGPT: A language model by OpenAI. <https://chat.openai.com/>, 2024. Accessed: 2024-09-17.
- [Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 311–318, 2002.
- [Pei *et al.*, 2023] Qizhi Pei, Wei Zhang, Jinhua Zhu, Kehan Wu, Kaiyuan Gao, Lijun Wu, Yingce Xia, and Rui Yan. Biot5: Enriching cross-modal integration in biology with chemical knowledge and natural language associations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1102–1123, 2023.
- [Pei *et al.*, 2024] Qizhi Pei, Lijun Wu, Kaiyuan Gao, Xiaozhuan Liang, Yin Fang, Jinhua Zhu, Shufang Xie, Tao Qin, and Rui Yan. BioT5+: Towards generalized biological understanding with IUPAC integration and multi-task

- tuning. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1216–1240, 2024.
- [Rogers and Hahn, 2010] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010.
- [Schneider *et al.*, 2015] Nadine Schneider, Roger A Sayle, and Gregory A Landrum. 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):2111–2120, 2015.
- [Schneider *et al.*, 2016] Nadine Schneider, Nikolaus Stiefl, and Gregory A Landrum. What’s what: The (nearly) definitive guide to reaction role assignment. *Journal of chemical information and modeling*, 56(12):2336–2346, 2016.
- [Schwaller *et al.*, 2019] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, 2019.
- [Seo *et al.*, 2021] Seung-Woo Seo, You Young Song, June Yong Yang, Seohui Bae, Hankook Lee, Jinwoo Shin, Sung Ju Hwang, and Eunho Yang. Gta: Graph truncated attention for retrosynthesis. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, volume 35, pages 531–539, 2021.
- [Taori *et al.*, 2023] Rohan Taori, Ishaan Jain, Shubhu Lohia, Adrià Garriga-Alonso, Daphne Ippolito, et al. Stanford Alpaca: An instruction-following LLaMA model, 2023.
- [Taylor *et al.*, 2022] Ross Taylor, Marcin Kardas, Guillem Cucurull, et al. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*, 2022.
- [Touvron *et al.*, 2023] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and efficient foundation language models, 2023.
- [Tu and Coley, 2022] Zhengkai Tu and Connor W Coley. Permutation invariant graph-to-sequence model for template-free retrosynthesis and reaction prediction. *Journal of chemical information and modeling*, 62(15):3503–3513, 2022.
- [Ucak *et al.*, 2022] Umit V Ucak, Islambek Ashyrmamatov, Junsu Ko, and Juyong Lee. Retrosynthetic reaction pathway prediction through neural machine translation of atomic environments. *Nature Communications*, 13(1):1186, 2022.
- [van der Maaten and Hinton, 2008] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.
- [Wan *et al.*, 2022] Yue Wan, Chang-Yu Hsieh, Ben Liao, and Shengyu Zhang. Retroformer: Pushing the limits of end-to-end retrosynthesis transformer. In *Proceeding of the 39th International Conference on Machine Learning*, pages 22475–22490, 2022.
- [Wang *et al.*, 2023] Yu Wang, Chao Pang, Yuzhe Wang, Junru Jin, Jingjie Zhang, Xiangxiang Zeng, Ran Su, Quan Zou, and Leyi Wei. Retrosynthesis prediction with an interpretable deep-learning framework based on molecular assembly tasks. *Nature Communications*, 14(1):6155, 2023.
- [Xu *et al.*, 2023] Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6268–6278, 2023.
- [Zdravil *et al.*, 2024] Barbara Zdravil, Eloy Felix, Fiona Hunter, Emma J Manners, James Blackshaw, Sybilla Corbett, Marleen de Veij, Harris Ioannidis, David Mendez Lopez, Juan F Mosquera, et al. The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic Acids Research*, 52(D1):D1180–D1192, 2024.
- [Zhao *et al.*, 2025] Zihan Zhao, Da Ma, Lu Chen, Liangtai Sun, Zihao Li, Yi Xia, Bo Chen, Hongshen Xu, Zichen Zhu, Su Zhu, Shuai Fan, Guodong Shen, Kai Yu, and Xin Chen. Developing chemdfm as a large language foundation model for chemistry. *Cell Reports Physical Science*, 6(4):102523, 2025.