
ChemDFM: A Large Language Foundation Model for Chemistry

Zihan Zhao^{1*} Da Ma^{1*} Lu Chen^{1,2†} Liangtai Sun¹ Zihao Li³ Yi Xia²
Bo Chen² Hongshen Xu¹ Zichen Zhu¹ Su Zhu⁴ Shuai Fan⁴ Guodong Shen²
Kai Yu^{1,2†} Xin Chen^{2†}

¹X-LANCE Lab, Department of Computer Science and Engineering
MoE Key Lab of Artificial Intelligence, SJTU AI Institute
Shanghai Jiao Tong University, Shanghai, China

²Suzhou Laboratory, Suzhou, China

³Shanghai Key Laboratory for Molecular Engineering of Chiral Drugs
School of Chemistry and Chemical Engineering
Shanghai Jiao Tong University, Shanghai, China

⁴AI Speech Co., Ltd., Suzhou, China

{zhao_mengxin, chenlusz, kai.yu}@sjtu.edu.cn

Abstract

Artificial intelligence (AI) has played an increasingly important role in chemical research. However, most models currently used in chemistry are specialist models that require training and tuning for specific tasks. A more generic and efficient solution would be an AI model that could address many tasks and support free-form dialogue in the broad field of chemistry. In its utmost form, such a generalist AI chemist could be referred to as Chemical General Intelligence. Large language models (LLMs) have recently logged tremendous success in the general domain of natural language processing, showing emerging task generalization and free-form dialogue capabilities. However, domain knowledge of chemistry is largely missing when training general-domain LLMs. The lack of such knowledge greatly hinders the performance of generalist LLMs in the field of chemistry. To this end, we develop ChemDFM, a pioneering LLM for chemistry trained on 34B tokens from chemical literature and textbooks, and fine-tuned using 2.7M instructions. As a result, it can understand and reason with chemical knowledge in free-form dialogue. Quantitative evaluations show that ChemDFM significantly surpasses most representative open-source LLMs. It outperforms GPT-4 on a great portion of chemical tasks, despite the substantial size difference. We have open-sourced the inference codes, evaluation datasets, and model weights of ChemDFM on Huggingface³.

1 Introduction

With the rapid development of artificial intelligence (AI), utilizing AI to assist chemical research has garnered increasing attention [Wang et al., 2023b, Back et al., 2024]. Various AI models have been developed for tasks such as property prediction [Zhou et al., 2022, Wu et al., 2023b, Chen et al., 2023], molecular captioning and generation [Xu et al., 2021, Edwards et al., 2022, Perron et al., 2022,

*Zihan Zhao and Da Ma contribute equally to this work.

†Xin Chen, Kai Yu, and Lu Chen are the corresponding authors.

³<https://huggingface.co/OpenDFM/ChemDFM-v1.0-13B>

Du et al., 2024, Lu et al., 2024], and reaction predictions [Schwaller et al., 2020, Wang et al., 2021, Han et al., 2024]. Since BERT [Devlin et al., 2019] and GPT [Radford et al.], efforts have been made to fine-tune pre-trained models for specific chemical tasks [Zhou et al., 2022, Edwards et al., 2022, Liu et al., 2023, Luo et al., 2023, Zhang et al., 2024]. However, these models are typically trained on a meticulously curated dataset to solve a designated task in a particular scenario, leading to a one-to-one relationship between models and tasks. Once out of that specific scenario, they are often not useful, even for highly related tasks. A more attractive and practical AI system should be capable of handling a wide range of chemical tasks under real-world scenarios and conducting free-form human-AI collaborations. Such an AI system necessitates a comprehensive array of chemical competencies, coupled with the ability to comprehend and reason in both chemical and natural languages. This would enable it to work as a research assistant or even collaborator alongside human researchers. This could be an essential step towards eventually achieving Chemical Artificial General Intelligence.

In pursuit of a highly integrated AI system for a broad range of chemical challenges, recent advancements in large language models (LLMs) [Du et al., 2022, Touvron et al., 2023a, Xu et al., 2023] brought great new hopes. Numerous studies have demonstrated the remarkable competencies of LLMs in natural language understanding and task generalization [Wei et al., 2021, Xu et al., 2023], deductive reasoning [Wei et al., 2022, Kojima et al., 2022], and tool utilization [Schick et al., 2023, Qin et al., 2024]. These made LLMs shine in traditional natural language processing tasks and accomplish problems that were previously unimaginable and unsolvable, such as handling tasks in unknown scenarios or conducting free-form dialogues with humans. These inherent strengths underscore the viability of employing LLMs as AI-driven research collaborators in the field of chemistry.

Different from general domains, tasks in chemical domains necessitate models to possess additional chemical comprehension capabilities for understanding and reasoning over chemical-specialized language and knowledge. This hinders general domain LLMs from excelling in chemical tasks as they often lack in-depth chemical knowledge [Kristiadi et al., 2024]. For example, molecules are a vital component of the chemical world. Although molecules can be conveyed through natural-language-like notations such as SMILES (Simplified Molecular Input Line Entry System), IUPAC names, and molecular formulas, their meanings and intrinsic structures are entirely different from those in natural language. CO represents carbon monoxide in chemistry, not Colorado, while Co represents Cobalt, not a company, and (CO) as part of a SMILES typically represents the carbonyl group. The lack of understanding of these molecular notations severely limits the applicability and performance of general domain LLMs in solving chemistry problems. Therefore, we believe that equipping general-domain LLMs with rich chemical knowledge of task-specific chemical models, as illustrated in Figure 1, is vital for developing LLMs useful in the field of chemistry.

In this work, we propose **ChemDFM**, a Dialogue Foundation Model for Chemistry. ChemDFM takes advantage of the pre-trained LLaMa-13B model [Touvron et al., 2023a], an open-source general-domain LLM, and is further specialized in chemistry through two phases: 1) **Domain Pre-training**, where the model harvests the chemical knowledge from research articles and textbooks, and 2) **Instruction Tuning**, where the model familiarizes itself with chemical language and patterns, especially molecule notations. Each phase uses an extensive and diverse collection of chemical data: 1) nearly 34B tokens from over 3.8M chemical papers and 1.4K textbooks in chemistry used in Phase I, and 2) over 2.7M instructions crafted from various chemical databases in Phase II. Apart from chemical data, we also incorporated a substantial amount of general-domain data in both phases to make sure that ChemDFM maintains comprehension and reasoning capabilities of natural language while acquiring new chemical knowledge. As a result, ChemDFM can simultaneously handle a wide range of chemical tasks and convey free-form dialogues using the language of chemists, enabling human-AI collaboration in chemical research.

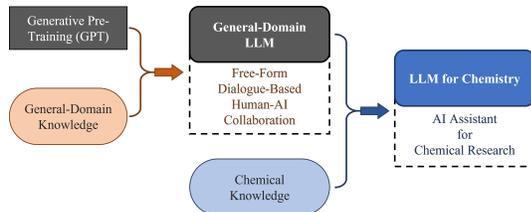


Figure 1: Scheme to obtain an LLM for chemistry, through using chemical domain knowledge to train a general-domain LLM.

A series of experiments have been conducted to evaluate the prowess of ChemDFM, including molecule recognition, molecule design, molecular property prediction, and reaction analysis. The results show that ChemDFM achieves advanced performances, surpassing typical open-source LLMs.

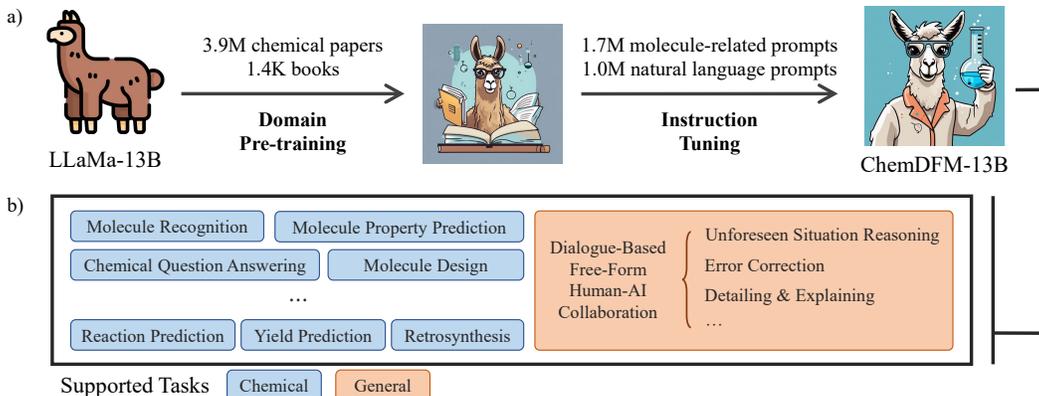


Figure 2: a) Two-step training procedure to obtain ChemDFM. The icons are generated by the SDXL model provided by Stability AI⁵. b) Various types of tasks ChemDFM is capable of.

It even outperforms GPT-4 on most tasks, despite the notable difference in model size. We further compared the performance between ChemDFM and the baseline LLMs in free-form unseen scenarios analogous to real-world scenarios. The test examples were constructed based on the latest chemical papers to avoid possible data leakage. The results show that ChemDFM can generate answers that are more accurate and relevant to the specific questions. These findings suggest that ChemDFM, capable of handling a broad range of chemical tasks and reasoning in both chemical and natural languages, can indeed serve as an AI assistant in chemical research.

2 ChemDFM

As outlined in Figure 2, ChemDFM is trained based on LLaMa, a general domain LLM. Domain knowledge of chemistry is instilled in ChemDFM in two steps: Domain Pre-training and Instruction Tuning. Through this two-stage specialization process, ChemDFM “learned” chemistry and gained abilities such as molecule recognition and reaction prediction. The training process is presented below and evaluations of ChemDFM’s capability are elaborated in the next section.

2.1 Domain Pre-training

Data used to train general-domain LLMs must contain knowledge covering a wide range of topics. Such broadness is often accompanied by sacrifices of deepness in each field. While models trained on such data have successfully gained strong natural language understanding and reasoning capabilities, they often fall short when it comes to in-depth specialized knowledge. The lack of domain knowledge is partially responsible for the well-known “hallucination” problem [Huang et al., 2023]. To alleviate this problem, we collected a corpus of data rich in chemical knowledge for domain pre-training, primarily from the two most authoritative sources for chemical knowledge: textbooks and published papers. Textbooks represent the widely accepted knowledge and basic principles of chemistry while published papers offer more details and more up-to-date chemical knowledge, some of which have not been incorporated into textbooks. Specifically, we selected 1.4K chemistry books from LibreTexts⁶ and Gold Books⁷ and collected 3.9M open-access papers in chemistry-related topics before January 2022. After further pre-processing and deduplication, we obtained 49M tokens from the textbooks and 34B tokens from the published research articles. To maintain the LLM’s general-domain knowledge and capabilities, we also included highly selective data in the general field, including Wikipedia, Arxiv, Books, StackExchange, GitHub code, WuDao Corpora [Yuan et al., 2021], etc. More details of domain pre-training are available in Appendix Section A.1.

⁵<https://stability.ai/>

⁶<https://libretexts.org/>

⁷<https://goldbook.iupac.org/>

2.2 Instruction Tuning

The data for the chemical instruction tuning dataset comprises two main categories: chemical knowledge presented in natural language and specialized molecular notations. A dataset containing over 1M question-answering pairs specialized in chemistry was constructed for instruction tuning to enhance the model’s capability to process chemistry-related natural language. These data were generated from two sources. The first one is established question-answering datasets, including ARC [Clark et al., 2018], PIQA [Bisk et al., 2020], SciQ [Welbl et al., 2017], and HendrycksTest [Hendrycks et al., 2021]. The other source of questions is middle school exams. We collected open-source exam questions from the Internet and constructed question-answer pairs (with key points or problem-solving strategies when available) for the instruction tuning of ChemDFM.

While natural languages such as English or German are generally descriptive and highly versatile, they are often not the best media to convey chemical knowledge. For example, it is often much easier and more comprehensible to draw the molecular structure of a complicated organic molecule than to describe it using natural language. Generations of chemists have developed many specialized notations, such as molecular formulas and Simplified Molecular Input Line Entry System (SMILES) [Weininger, 1988] notation. This represents a key challenge for LLMs to understand chemistry. A key goal of the instruction tuning stage was to tackle this challenge by familiarizing ChemDFM with the specialized notations. In training ChemDFM, we chose SMILES, one of the most popular notations of molecules, as the main representation for molecules. It uses a sequence of letters to present a molecule, retaining rich structural information such as molecular configuration and chirality in most cases. In addition, its text-like data structure makes it highly compatible with LLMs.

To help the model comprehend SMILES, three kinds of molecular data were used: 1. **Molecule description (MD) and text-based molecule design (TBMD)**. Our dataset includes all the SMILES-description pairs from PubChem⁸, a web-scale chemical database that contains more than 100M compounds. The model was instructed to generate descriptions of given molecules or reversely, generate molecule(s) that match a description. We duplicated samples with descriptions longer than two sentences to further enhance the data quality. 2. **Molecular property prediction (MPP)**. The model was instructed to predict the properties of a given molecule. This data was constructed based on the widely used molecular property prediction benchmark, MoleculeNet [Wu et al., 2018]. 3. **Reaction completion (RC)**. The model was also instructed to complete chemical reactions in which one or more reactants/products were masked randomly. The reactions were sampled from USPTO [Lowe, 2012], the largest open-source chemical reaction database.

In addition to SMILES, we indirectly include two other widely used notations of molecules, IUPAC names and molecular formulas. We instructed the model to translate between the three notations, e.g. predicting SMILES of a molecule

Table 1: Itemized list of our instruction tuning dataset. MD: Molecule Description, TBMD: Text-Based Molecule Design, MPP: Molecular Property Prediction, RC: Reaction Completion, MNA: Molecular Notation Alignment.

Data Type	# samples	Data Source
QAs from Datasets	131K	ARC, PIQA, SciQ HendrycksTest
QAs from Exam	915K	Internet
MD	576K	PubChem
TBMD	576K	PubChem
MPP	102K	MoleculeNet
RC	300K	USPTO
MNA	120K	PubChem

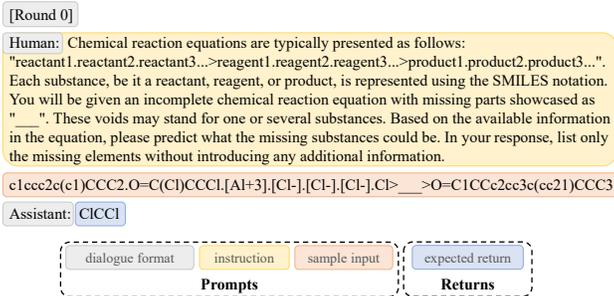


Figure 3: Representative question used in instruction tuning.

⁸<https://pubchem.ncbi.nlm.nih.gov/>

given its IUPAC name and vice versa, allowing it to understand these alternatives. This kind of data is called **Molecular Notation Alignment (MNA)** in this work.

Table 1 lists the itemized entries of our instruction tuning dataset. All the data samples take the form of (prompt, returns) tuples, where the prompt is composed of the dialogue format, instructions, and example inputs, and the returns are the expected outputs. Such an example is presented in Figure 3. To diversify the natural language instructions, we used GPT-4 to rephrase instructions for all tasks. The number of different instructions for each task ranges from 20 to 200.

To maintain the advanced natural language comprehension abilities, we also included a substantial amount of general domain data for the instruction-tuning of ChemDFM. The ratio of data from the chemical domain to the general domain is roughly 1:2. The instruction tuning of ChemDFM is a full-parameter tuning process with more details in Appendix Section A.2.

3 Evaluations

To assess ChemDFM’s capability in chemistry, we compared its performance against several generalist LLM models: GPT-4 [OpenAI, 2023], LLaMa-2 [Touvron et al., 2023b] and Galactica [Taylor et al., 2022], as they represent very large generalist LLMs, medium-sized generalist LLMs and LLMs tuned for science, respectively. We used ChemLLMBench [Guo et al., 2023] for quantitative evaluation of ChemDFM’s ability in chemistry and then carried out qualitative analyses of ChemDFM’s free-form collaboration capacity, focusing on its superior chemistry-related conversation processing power.

3.1 Quantitative Evaluation

ChemLLMBench is made of a series of chemical tasks covering a wide range of chemistry-related topics. The standard form of evaluation was conducted on 100 instances randomly sampled from the respective test sets of the tasks. To ensure a fair comparison, we used the same 100 samples when comparing different LLMs, unless otherwise specified. Some non-LLM task-specific models were used for comparisons when available. Detailed explanations of the tasks and the prompt format for ChemDFM can be found in Appendix Section B. Specifically, the quantitative evaluation tasks can be categorized into the following four groups.

1) Molecule recognition. There are two series of tasks in ChemLLM-Bench that directly assess the capability to recognize molecules: *name prediction* and *molecule captioning*. In the *name prediction* tasks, a model is asked to translate between different notations for molecules, including SMILES, IUPAC name, and molecular formula. Specifically, it consists of four tasks: SMILES to IUPAC name translation (S2I), IUPAC name to SMILES translation (I2S), SMILES to Molecular Formula translation (S2MF), and IUPAC name to Molecular Formula translation (I2MF). For IUPAC names and SMILES, we normalized the predictions before calculating the accuracy scores, while for molecular formulas, only exact matches are considered correct answers. The *molecule captioning* tasks further require LLMs to not only recognize the molecule present by a given SMILES notation but also generate a brief description of it using natural language. In these tasks, traditional captioning metrics like BLUE, ROUGE, and METEOR are used to assess the model’s performance on a test set of ChEBI-20 [Edwards et al., 2021].

Benchmark performance of different models on these two molecule recognition tasks is reported in Table 2 and Table 3, respectively. Table 2 shows that most LLMs, including GPT-4, can hardly complete *name prediction* tasks, indicating a limited understanding of molecules and ChemDFM outperforms open-source LLMs by a significant margin across all these tasks. This outstanding

Table 2: Accuracy scores in name prediction tasks. Baseline results are from Guo et al. [2023]. S2I: SMILES to IUPAC names translation, I2S: IUPAC names to SMILES translation, S2MF: SMILES to molecule formulas translation, I2MF: IUPAC names to molecule formulas translation.

Model	S2I↑	I2S↑	S2MF↑	I2MF↑
<i>Task-specific specialist models</i>				
STOUT	55.0	70.0	-	-
<i>LLM-based generalist models</i>				
GPT-4	0	1.2	8.6	8.4
LLaMa2-13B-chat	0	0	1.0	0
Galactica-30B	0	0	0	0
ChemDFM-13B	4.0	11.0	73.0	51.0

performance of different models on these two molecule recognition tasks is reported in Table 2 and Table 3, respectively. Table 2 shows that most LLMs, including GPT-4, can hardly complete *name prediction* tasks, indicating a limited understanding of molecules and ChemDFM outperforms open-source LLMs by a significant margin across all these tasks. This outstanding

Table 3: Benchmark results of different models in molecule captioning tasks. †: results from Guo et al. [2023]. *: reproduced results.

Model	BLEU-2↑	BLEU-4↑	ROUGE-1↑	ROUGE-2↑	ROUGE-L↑	METEOR↑
<i>Task-specific specialist models</i>						
Text+Chem T5 [Christofidellis et al., 2023]	0.625	0.542	0.682	0.543	0.622	0.648
MolXPT [Liu et al., 2023]	0.594	0.505	0.660	0.511	0.597	0.626
InstructMol [Cao et al., 2023a]	0.475	0.371	0.566	0.394	0.502	0.509
Mol-Instruction [Fang et al., 2023]	0.249	0.171	0.331	0.203	0.289	0.271
<i>LLM-based generalist models</i>						
GPT-4 (10-shot)†	0.464	0.365	0.545	0.362	0.459	0.519
GPT-4 (0-shot)†	0.062	0.013	0.192	0.040	0.125	0.209
LLaMa-2-13B-chat (10-shot)†	0.197	0.140	0.331	0.193	0.265	0.372
Galactica-30B (10-shot)*	0.114	0.055	0.334	0.189	0.330	0.187
Galactica-30B (0-shot)†	0.008	0.002	0.019	0.004	0.015	0.043
ChemDFM-13B (0-shot)	<u>0.321</u>	<u>0.265</u>	<u>0.490</u>	0.374	0.483	<u>0.402</u>

Table 4: Benchmark results of different models in text-based molecule design tasks. †: results from Guo et al. [2023]. *: 10-shot results

Model	Exact↑	BLEU↑	Dis↓	Validity↑	MACCS↑	RDKit↑	Morgan↑
<i>Task-specific specialist models</i>							
MolXPT [Liu et al., 2023]	21.5	-	-	98.3	0.859	0.757	0.667
Text+Chem T5 [Christofidellis et al., 2023]	32.2	0.853	16.87	94.3	0.901	0.816	0.757
Mol-Instruction [Fang et al., 2023]	0.2	0.345	41.4	100	0.412	0.231	0.147
<i>LLM-based generalist models</i>							
GPT-4†*	17.4	0.816	21.2	88.8	0.867	0.738	0.672
LLaMa-2-13B-chat†*	2.0	0.626	34.0	78.2	0.679	0.568	0.454
Galactica-30B†	0.0	0.004	2738	95.6	0.233	0.109	0.053
ChemDFM-13B	45.0	0.874	9.9	98.0	0.922	0.871	0.798

performance of ChemDFM proves its robust molecule recognition capabilities and validates the effectiveness of our specialization process.

In *molecule captioning* tasks (as shown in Table 3), ChemDFM also performs far superior to open-source LLMs. The results denote that ChemDFM not only recognizes molecules but also infers their underlying chemical essence and nature. It is worth noting the drastic drop in GPT-4’s performance from the ten-shot setting to the zero-shot setting, which indicates that GPT-4 thrives mostly on its extraordinary natural language capabilities to learn from given exemplars while its inherent molecule recognition capability is relatively fragile. Comparatively, ChemDFM achieves comparable performance without the help of exemplars, demonstrating its intrinsic molecule recognition capability.

2) Text-based molecule design. To evaluate LLM’s efficiency in making qualified molecule designs, ChemLLMBench reverses the molecule captioning tasks and asks the models to generate molecules based on their descriptions. Specifically, in the *text-based molecule design* task, models are asked to predict the SMILES of the molecule that fits the given description. Two sets of metrics are utilized to measure the performance of these tasks. The first set measures the text-based similarity of the predicted SMILES compared to the golden SMILES, which includes exact match, BLUE, and Levenshtein distance. The second set of metrics measures the chemical similarity of the predicted molecules to the golden molecules, including the validity of the predicted SMILES and the FTS (fingerprint Tanimoto Similarity) [Tanimoto, 1958] in terms of MACCS [und David Metzener, 1988], RDKit⁹, Morgan [Morgan, 1965].

As shown in Table 4, ChemDFM outperforms not only the generalist LLMs but also the traditional task-specific specialist models across almost all metrics, which is both surprising and promising. Considering that task-specific specialist models were evaluated on the entire test set, whereas the performance of ChemDFM was initially assessed on only 100 samples, we further evaluated ChemDFM on the complete test set to align with the task-specific models for a fair comparison. The results, shown in Table 7 of the Appendix, further validate the advantage of ChemDFM. The results from Table 4 and 7 unveil two key superiorities of ChemDFM over other models. On the one hand, ChemDFM has effectively established a relationship between SMILES notations and the chemical

⁹<https://www.rdkit.org/>

Table 5: AUC-ROC scores [Bradley, 1997] of different models in molecular property prediction tasks. Avg: average. †: reproduced results (The results of GPT-4 were obtained in January 2024).

Model	BACE↑	BBBP↑	ClinTox↑	HIV↑	Tox21↑	Avg↑
<i>Task-specific specialist models</i>						
Uni-Mol [Zhou et al., 2022]	85.7	72.9	91.9	80.8	79.6	82.2
MolXPT [Liu et al., 2023]	88.4	80.0	95.3	78.1	77.1	83.8
InstructMol [Cao et al., 2023a]	85.9	64.0	-	74.0	-	-
<i>LLM-based generalist models</i>						
GPT-4†	62.5	61.5	51.6	65.9	55.2	59.3
LLaMa-2-13B-chat†	26.0	60.3	45.7	29.0	51.7	42.5
Galactica-30B [Taylor et al., 2022]	72.7	59.6	82.2	75.9	68.5	71.8
ChemDFM-13B	78.4	66.7	89.9	73.6	79.8	77.7

Table 6: Accuracy scores of different models in reaction prediction and retrosynthesis tasks. B-H: Buchwald-Hartwig dataset [Ahneman et al., 2018]. Suzuki: Suzuki-Miyaura dataset [Reizman et al., 2016]. YP: Yield Prediction, RP: Reactant Prediction, RS: Reagent Selection, Retro: Retrosynthesis. †: results from Guo et al. [2023]. Please refer to Table 9~12 in the Appendix for complete results.

Model	YP↑	RP↑	Retro↑	RS↑
<i>task-specific specialist models</i>				
UAGNN [Kwon et al., 2022]	96.1	-	-	-
Chemformer [Irwin et al., 2022]	-	93.8	53.6	-
<i>LLM-based generalist models</i>				
GPT-4†	<u>78.2</u>	<u>23.0</u>	<u>11.4</u>	45.3
LLaMa-2-13B-chat†	0.7	3.2	0.0	16.0
Galactica (30B)†	0.4	3.6	1.6	8.0
ChemDFM-13B	81.0	49.0	12.0	<u>23.7</u>

nature of compounds in our model, which other LLMs lack. On the other hand, ChemDFM benefits from the solid natural language comprehension capabilities inherited from LLaMa, which task-specific specialist models lack. Altogether, ChemDFM constructs a more comprehensive knowledge system in chemistry, which helps it surpass both generalist and task-specific specialist models.

3) Molecular property prediction. The *molecular property prediction* tasks in ChemLLMBench consist of five tasks from the MoleculeNet [Wu et al., 2018], including BACE, BBBP, HIV, ClinTox, and Tox21. Among them, BACE and BBBP each contain a single balanced binary classification task. HIV contains a single unbalanced binary classification task. ClinTox and Tox21 comprise two and twenty-one unbalanced binary classification tasks, respectively. To address the severe label imbalance in these tasks, the Area Under the Curve of the Receiver Operating Characteristic (AUC-ROC) metric [Bradley, 1997] was introduced. To better assess the molecular property prediction, we adopted a scaffold-vertical manner for data splitting. Specifically, the molecules from the DeepChem library [Ramsundar et al., 2019] were first grouped based on their Bemis-Murcko scaffold [Bemis and Murcko, 1996] representations. The datasets were then split into training and test sets according to these groups. This method ensures that no molecule sharing the same scaffold would appear in both the training set and the test set. While avoiding information leaking due to mere similarity of molecules, this method also significantly increases the difficulty of the tasks, making the assessment more challenging and meaningful. The results listed in Table 5 show that ChemDFM consistently outperforms other LLMs in all but one molecular property prediction task.

4) Reaction prediction and retrosynthesis. ChemLLMBench includes four types of tasks targeted at evaluating models’ capability of reaction understanding: *Yield Prediction* (YP), *Reaction Prediction* (RP), *Reagent Selection* (RS), and *Retrosynthesis* (Retro). The *yield prediction* tasks ask models to predict whether the given reaction is a high-yield reaction and are constructed based on two High-Throughput experimentation (HTE) datasets: the Buchwald-Hartwig dataset [Ahneman et al., 2018] and the Suzuki-Miyaura dataset [Reizman et al., 2016]. The *reaction prediction* tasks ask models to predict the product of the given reaction. The USPTO-MIT dataset [Jin et al., 2017]

LLM-based reading partner or assistant can provide instant explanations and answers to such questions. In Figure 4, we compare the answers generated by ChemDFM with those from other LLMs. We have provided three example questions, with more examples in Appendix Section C.1, which are generally consistent with the analysis below. To prevent information leakage, the questions were constructed from chemistry papers published in 2023 only. Since ChemDFM only learned from papers published before 2022, this approach ensures ChemDFM has not learned the answers during training and simulates ChemDFM’s performance as a reading partner or tutor when reading new papers. **Q1** represents a question of widely known domain knowledge. All LLMs including ChemDFM provide good answers. However, when questions involve new molecules and reactions (**Q2** [Yin et al., 2023] & **Q3** [Dargo et al., 2023]), the performances differ. Specifically, LLaMa-2 and Galactica primarily rely on retrieving knowledge from memory, which can result in numerous knowledge points that are correct but irrelevant or even misleading in the context of the questions. GPT-4 shows a primary level of ability to answer questions based on the provided molecules and/or reactions. It effectively answers **Q2** but struggles with more complex questions involving complicated molecules such as **Q3**. In **Q3**, GPT-4 fails to fully recognize the underlying chemical aspects of the question and proposes methods that could violate the molecule’s catalytic activity. It is also worth noticing that as GPT-4 is a closed-source LLM, it is uncertain whether the literature used to construct the questions is included in GPT-4’s training corpus. Therefore, these "new papers" may not be new to GPT-4. In contrast to other LLMs, ChemDFM shows the ability to integrate memory-based knowledge while considering the situation described in the questions, providing key points that are highly relevant to the question. In terms of accuracy, relevance, and overall quality of the answers, ChemDFM largely outperforms other LLMs including GPT-4, demonstrating a better understanding of molecules and reactions, especially in the example of **Q3**. Apart from presenting key points, ChemDFM also endeavors to expand on its explanation and elaborate on the mechanism of the queried reactions or the proposed solutions, although this occasionally leads to inaccurate answers, as seen in the cases of **Q1** and **Q2**. Please refer to Appendix Section C.1 for a more detailed analysis.

A knowledgeable discussion partner who is always available and patient would be invaluable for researchers, particularly in the stage of hypothesis generation and design of experiment (DOE). Figure 5 illustrates a scenario inspired by Yin et al. [2023] that showcases ChemDFM’s potential to assist researchers in free-form dialogues as an AI research partner. In this example, a human researcher aimed to selectively oxidize one of the two carbonyl groups of a molecule. The initial solution given by ChemDFM would lead to the oxidation of both carbonyl groups. However, after being alerted and challenged by the human researcher, ChemDFM acknowledged the mistake and proposed two possible strategies: using a weaker oxidation agent/condition or introducing a protecting group. Upon the researcher’s decision to use a protecting group, ChemDFM provided detailed recommendations, including a feasible agent and reaction condition. Throughout the dialogue, ChemDFM exhibited promising capabilities in comprehension (Round 1), error correction (Round 2), and detailing (Round 3), showcasing its efficacy in mastering both chemical and natural language. More examples can be found in Appendix Section C.2.

4 Related Work

There have been several pioneering studies focusing on leveraging LLMs to solve chemical problems. These works typically adopt one of two general strategies. The first one treats LLMs as powerful base models for multi-task training, neglecting their greatest strength in natural language understanding and reasoning [Christofidellis et al., 2023, Fang et al., 2023, Cao et al., 2023a, Zheng et al., 2023, Kim et al., 2024, Yu et al., 2024]. Consequently, the models devised under this framework are confined to solving the specific tasks on which they were trained, losing the ability to tackle unseen tasks or conduct free-form human-AI collaborations. The other strategy exploits LLMs’ strong natural language understanding and reasoning abilities, using them directly to handle complex chemical tasks described in natural language [Hatakeyama-Sato et al., 2023, Cao et al., 2023b, Boiko et al., 2023, Yoshikawa et al., 2023, M. Bran et al., 2024, Ruan et al., 2024]. However, most of them suffer from the fact that generalist LLMs lack an inherent understanding of chemical language and knowledge [Kristiadi et al., 2024]. We argue that an LLM useful in chemistry must learn and reason with both general-domain knowledge and chemical knowledge. In this work, we tried to achieve this by equipping general-domain LLMs with rich chemical knowledge of task-specific chemical models and obtained promising results.

Notably, this strategy has been successfully applied to develop specialist LLMs for several other scientific domains. For example, Med-PaLM [Singhal et al., 2023] and PMC-LLaMa [Wu et al., 2023a] are specialized LLMs for biology and medicine. Similarly, ChatDoctor [Li et al., 2023] and DrugChat [Liang et al., 2023] also offer LLMs specifically for the medicine field but focus on medical inquiries and drug discoveries. Other domain-specific LLMs have endeavored include education [Dan et al., 2023], materials science [Xie et al., 2023], and geography [Deng et al., 2023]. It is worth noting that most of these works only focus on natural language. Domain-specific languages, which differ significantly from natural languages, such as SMILES in chemistry, are often overlooked.

5 Conclusion

In summary, this paper introduces ChemDFM, a specialist LLM that evolves from a generalist LLM through pre-training and instruction tuning using domain knowledge in chemistry. Quantitative evaluations show ChemDFM’s strong comprehension of molecular notations and reasoning capabilities for chemical knowledge, resulting in excellent performance in a wide range of chemical tasks such as molecular design and reaction analysis. In scenarios such as paper reading and experimental design, ChemDFM shows great potential in wielding chemical and natural languages to assist researchers through dialogue-based, free-form human-AI collaborations.

Acknowledgments and Disclosure of Funding

This work was supported by the National Science and Technology Major Project 2023ZD0120703, the China NSFC Projects (U23B2057, 62106142 and 62120106006), and the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102).

References

- Derek T Ahneman, Jesús G Estrada, Shishi Lin, Spencer D Dreher, and Abigail G Doyle. Predicting reaction performance in c–n cross-coupling using machine learning. *Science*, 360(6385):186–190, 2018.
- Seoin Back, Alán Aspuru-Guzik, Michele Ceriotti, Ganna Gryn’ova, Bartosz Grzybowski, Geun Ho Gu, Jason Hein, Kedar Hippalgaonkar, Rodrigo Hormázabal, Yousung Jung, et al. Accelerated chemical science with ai. *Digital Discovery*, 3(1):23–33, 2024.
- Guy W Bemis and Mark A Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of medicinal chemistry*, 39(15):2887–2893, 1996.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI*, volume 34, pages 7432–7439, 2020.
- Daniil A Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- Andrew P Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.
- 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. *arXiv preprint arXiv:2311.16208*, 2023a.
- Zhonglin Cao, Rishikesh Magar, Yuyang Wang, and Amir Barati Farimani. Moformer: self-supervised transformer model for metal–organic framework property prediction. *Journal of the American Chemical Society*, 145(5):2958–2967, 2023b.
- Marta Catellani, Franco Frignani, and Armando Rangoni. A complex catalytic cycle leading to a regioselective synthesis of o, o’-disubstituted vinylarenes. *ChemInform*, 28(16), 1997.

- Tianlang Chen, Shengjie Luo, Di He, Shuxin Zheng, Tie-Yan Liu, and Liwei Wang. GeoMFormer: A general architecture for geometric molecular representation learning. In *NeurIPS 2023 AI for Science Workshop*, 2023. URL <https://openreview.net/forum?id=s0UNTuuqU5>.
- Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo Manica. Unifying molecular and textual representations via multi-task language modelling. In *ICML*, 2023.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Yuhao Dan, Zhikai Lei, Yiyang Gu, Yong Li, Jianghao Yin, Jiaju Lin, Linhao Ye, Zhiyan Tie, Yougen Zhou, Yilei Wang, Aimin Zhou, Ze Zhou, Qin Chen, Jie Zhou, Liang He, and Xipeng Qiu. Educhat: A large-scale language model-based chatbot system for intelligent education. *arXiv preprint arXiv:2308.02773*, 2023.
- Gyula Dargo, David Kis, Martin Gede, Sushil Kumar, Jozsef Kupai, and Gyorgy Szekely. Mesesamol, a bio-based and versatile polar aprotic solvent for organic synthesis and depolymerization. *Chemical Engineering Journal*, page 144365, 2023.
- Cheng Deng, Tianhang Zhang, Zhongmou He, Yi Xu, Qiyuan Chen, Yuanyuan Shi, Luoyi Fu, Weinan Zhang, Xinbing Wang, Chenghu Zhou, Zhouhan Lin, and Junxian He. K2: A foundation language model for geoscience knowledge understanding and utilization. *arXiv preprint arXiv:2306.05064*, 2023.
- DB Dess and JC Martin. Readily accessible 12-i-5 oxidant for the conversion of primary and secondary alcohols to aldehydes and ketones. *The Journal of Organic Chemistry*, 48(22):4155–4156, 1983.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the NAACL*, 2019. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Yuanqi Du, Arian R Jamasb, Jeff Guo, Tianfan Fu, Charles Harris, Yingheng Wang, Chenru Duan, Pietro Liò, Philippe Schwaller, and Tom L Blundell. Machine learning-aided generative molecular design. *Nature Machine Intelligence*, pages 1–16, 2024.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. GLM: General language model pretraining with autoregressive blank infilling. In *Proceedings of the ACL*, 2022. doi: 10.18653/v1/2022.acl-long.26. URL <https://aclanthology.org/2022.acl-long.26>.
- Carl Edwards, ChengXiang Zhai, and Heng Ji. Text2Mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021. doi: 10.18653/v1/2021.emnlp-main.47. URL <https://aclanthology.org/2021.emnlp-main.47>.
- Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. Translation between molecules and natural language. In *Proceedings of the EMNLP*, 2022. doi: 10.18653/v1/2022.emnlp-main.26. URL <https://aclanthology.org/2022.emnlp-main.26>.
- 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. *arXiv preprint arXiv:2306.08018*, 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. *arXiv preprint arXiv:2305.18365*, 2023.
- Yuqiang Han, Xiaoyang Xu, Chang-Yu Hsieh, Keyan Ding, Hongxia Xu, Renjun Xu, Tingjun Hou, Qiang Zhang, and Huajun Chen. Retrosynthesis prediction with an iterative string editing model. *Nature Communications*, 15(1):6404, 2024.

- Yu Hao, Zi-Hao Li, Zhi-Gang Ma, Ru-Xin Liu, Rui-Tian Ge, Quan-Zhe Li, Tong-Mei Ding, and Shu-Yu Zhang. Axially chiral styrene-based organocatalysts and their application in asymmetric cascade michael/cyclization reaction. *Chemical Science*, 14(35):9496–9502, 2023.
- Kan Hatakeyama-Sato, Naoki Yamane, Yasuhiko Igarashi, Yuta Nabae, and Teruaki Hayakawa. Prompt engineering of gpt-4 for chemical research: what can/cannot be done? *Science and Technology of Advanced Materials: Methods*, (1):2260300, 2023.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *ICLR*, 2021.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions, 2023. URL <https://arxiv.org/abs/2311.05232>.
- Ross Irwin, Spyridon Dimitriadis, Jiazhen He, and Esben Jannik Bjerrum. Chemformer: a pre-trained transformer for computational chemistry. *Machine Learning: Science and Technology*, 3(1): 015022, 2022.
- Wengong Jin, Connor Coley, Regina Barzilay, and Tommi Jaakkola. Predicting organic reaction outcomes with weisfeiler-lehman network. *Advances in neural information processing systems*, 30, 2017.
- Seongmin Kim, Yousung Jung, and Joshua Schrier. Large language models for inorganic synthesis predictions. *Journal of the American Chemical Society*, 2024.
- Milton C Kloetzel. The diels-alder reactions with maleic anhydride. *Org React*, 4:1–59, 1948.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *NeurIPS*, 35:22199–22213, 2022.
- Agustinus Kristiadi, Felix Strieth-Kalthoff, Marta Skreta, Pascal Poupart, Alán Aspuru-Guzik, and Geoff Pleiss. A sober look at llms for material discovery: Are they actually good for bayesian optimization over molecules? *arXiv preprint arXiv:2402.05015*, 2024.
- Youngchun Kwon, Dongseon Lee, Youn-Suk Choi, and Seokho Kang. Uncertainty-aware prediction of chemical reaction yields with graph neural networks. *Journal of Cheminformatics*, 14:1–10, 2022.
- Yunxiang Li, Zihan Li, Kai Zhang, Ruilong Dan, Steve Jiang, and You Zhang. Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. *Cureus*, 15(6), 2023.
- Youwei Liang, Ruiyi Zhang, Li Zhang, and Pengtao Xie. Drugchat: Towards enabling chatgpt-like capabilities on drug molecule graphs. *arXiv preprint arXiv:2309.03907*, 2023.
- Zequn Liu, Wei Zhang, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Ming Zhang, and Tie-Yan Liu. MolXPT: Wrapping molecules with text for generative pre-training. In *Proceedings of the ACL*, 2023. doi: 10.18653/v1/2023.acl-short.138. URL <https://aclanthology.org/2023.acl-short.138>.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Daniel Mark Lowe. *Extraction of chemical structures and reactions from the literature*. PhD thesis, University of Cambridge, 2012.
- Shuqi Lu, Lin Yao, Xi Chen, Hang Zheng, Di He, and Guolin Ke. 3d molecular generation via virtual dynamics. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=QvipGVdE6L>.

- Shengjie Luo, Tianlang Chen, Yixian Xu, Shuxin Zheng, Tie-Yan Liu, Liwei Wang, and Di He. One transformer can understand both 2d & 3d molecular data. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=vZTp1oPV3PC>.
- Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, pages 1–11, 2024.
- H. L. Morgan. The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. *Journal of Chemical Documentation*, 5(2):107–113, 1965. doi: 10.1021/c160017a018. URL <https://doi.org/10.1021/c160017a018>.
- OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Damith Perera, Joseph W Tucker, Shalini Brahmabhatt, Christopher J Helal, Ashley Chong, William Farrell, Paul Richardson, and Neal W Sach. A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. *Science*, 359(6374):429–434, 2018.
- Quentin Perron, Olivier Mirguet, Hamza Tajmouati, Adam Skiredj, Anne Rojas, Arnaud Gohier, Pierre Ducrot, Marie-Pierre Bourguignon, Patricia Sansilvestri-Morel, Nicolas Do Huu, et al. Deep generative models for ligand-based de novo design applied to multi-parametric optimization. *Journal of Computational Chemistry*, 43(10):692–703, 2022.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, dahai li, Zhiyuan Liu, and Maosong Sun. ToolLLM: Facilitating large language models to master 16000+ real-world APIs. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=dHng200Jjr>.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20. IEEE Press, 2020. ISBN 9781728199986.
- Bharath Ramsundar, Peter Eastman, Patrick Walters, Vijay Pande, Karl Leswing, and Zhenqin Wu. *Deep Learning for the Life Sciences*. O'Reilly Media, 2019. <https://www.amazon.com/Deep-Learning-Life-Sciences-Microscopy/dp/1492039837>.
- Brandon J Reizman, Yi-Ming Wang, Stephen L Buchwald, and Klavs F Jensen. Suzuki–miyaura cross-coupling optimization enabled by automated feedback. *Reaction chemistry & engineering*, 1(6):658–666, 2016.
- Yixiang Ruan, Chenyin Lu, Ning Xu, Jian Zhang, Jun Xuan, Jianzhang Pan, Qun Fang, Hanyu Gao, Xiaodong Shen, Ning Ye, et al. Accelerated end-to-end chemical synthesis development with large language models, 2024.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761*, 2023.
- 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.
- Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical science*, 11(12):3316–3325, 2020.

- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.
- Taffee T Tanimoto. Elementary mathematical theory of classification and prediction. *Journal of Biomedical Science and Engineering*, 1958.
- Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*, 2022.
- 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. *arXiv preprint arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- John W. Ratcliff und David Metzener. Pattern matching: The gestalt approach. *Dr. Dobb's Journal*, 1988.
- Ai-Fang Wang, Jin-Miao Tian, Xiao-Jing Zhao, Zi-Hao Li, Ye Zhang, Ka Lu, Hong Wang, Shu-Yu Zhang, Yong-Qiang Tu, Tong-Mei Ding, et al. Asymmetric intramolecular hydroalkylation of internal olefin with cycloalkanone to directly access polycyclic systems. *Angewandte Chemie International Edition*, 62(39):e202308858, 2023a.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023b.
- Xiaorui Wang, Yuquan Li, Jiezhong Qiu, Guangyong Chen, Huanxiang Liu, Benben Liao, Chang-Yu Hsieh, and Xiaojun Yao. Retroprime: A diverse, plausible and transformer-based method for single-step retrosynthesis predictions. *Chemical Engineering Journal*, 420:129845, 2021.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *ICLR*, 2021.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*, 35: 24824–24837, 2022.
- David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- Johannes Welbl, Nelson F Liu, and Matt Gardner. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*, 2017.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-llama: Towards building open-source language models for medicine. *arXiv preprint arXiv:2304.14454*, 2023a.
- Fang Wu, Dragomir Radev, and Stan Z Li. Molformer: Motif-based transformer on 3d heterogeneous molecular graphs. In *Proceedings of the AAAI*, volume 37, pages 5312–5320, 2023b.
- Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.
- Tong Xie, Yuwei Wan, Wei Huang, Zhenyu Yin, Yixuan Liu, Shaozhou Wang, Qingyuan Linghu, Chunyu Kit, Clara Grazian, Wenjie Zhang, Imran Razzak, and Bram Hoex. Darwin series: Domain specific large language models for natural science. *arXiv preprint arXiv:2308.13565*, 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 EMNLP*, 2023. doi: 10.18653/v1/2023.emnlp-main.385. URL <https://aclanthology.org/2023.emnlp-main.385>.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2021.
- Zhewei Yao, Reza Yazdani Aminabadi, Olatunji Ruwase, Samyam Rajbhandari, Xiaoxia Wu, Ammar Ahmad Awan, Jeff Rasley, Minjia Zhang, Conglong Li, Connor Holmes, Zhongzhu Zhou, Michael Wyatt, Molly Smith, Lev Kurilenko, Heyang Qin, Masahiro Tanaka, Shuai Che, Shuaiwen Leon Song, and Yuxiong He. DeepSpeed-Chat: Easy, Fast and Affordable RLHF Training of ChatGPT-like Models at All Scales. *arXiv preprint arXiv:2308.01320*, 2023.
- Jun-Jie Yin, Yun-Peng Wang, Jun Xue, Feng-Fan Zhou, Xing-Qian Shan, Rong Zhu, Kun Fang, Lei Shi, Shu-Yu Zhang, Si-Hua Hou, et al. Total syntheses of polycyclic diterpenes phomopsene, methyl phomopsenonate, and iso-phomopsene via reorganization of c-c single bonds. *Journal of the American Chemical Society*, 145(39):21170–21175, 2023.
- Naruki Yoshikawa, Marta Skreta, Kourosh Darvish, Sebastian Arellano-Rubach, Zhi Ji, Lasse Bjørn Kristensen, Andrew Zou Li, Yuchi Zhao, Haoping Xu, Artur Kuramshin, et al. Large language models for chemistry robotics. *Autonomous Robots*, 47(8):1057–1086, 2023.
- Botao Yu, Frazier N Baker, Ziqi Chen, Xia Ning, and Huan Sun. Llasmol: Advancing large language models for chemistry with a large-scale, comprehensive, high-quality instruction tuning dataset. *arXiv preprint arXiv:2402.09391*, 2024.
- Sha Yuan, Hanyu Zhao, Zhengxiao Du, Ming Ding, Xiao Liu, Yukuo Cen, Xu Zou, Zhilin Yang, and Jie Tang. Wudaocorpora: A super large-scale chinese corpora for pre-training language models. *AI Open*, 2:65–68, 2021.
- Juzheng Zhang, Yatao Bian, Yongqiang Chen, and Quanming Yao. Unimot: Unified molecule-text language model with discrete token representation. *arXiv preprint arXiv:2408.00863*, 2024.
- Zhiling Zheng, Ali H Alawadhi, Saumil Chheda, S Ephraim Neumann, Nakul Rampal, Shengchao Liu, Ha L Nguyen, Yen-hsu Lin, Zichao Rong, J Ilja Siepmann, et al. Shaping the water-harvesting behavior of metal-organic frameworks aided by fine-tuned gpt models. *Journal of the American Chemical Society*, 145(51):28284–28295, 2023.
- Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. In *ICLR*, 2022.
- Qing-Bo Zhuang, Jin-Rui Tian, Ka Lu, Xiao-Ming Zhang, Fu-Min Zhang, Yong-Qiang Tu, Rong Fan, Zhi-Hao Li, and Yu-Dong Zhang. Catalytic asymmetric polycyclization of tertiary enamides with silyl enol ethers: Total synthesis of (-)-cephalocyclidin a. *Journal of the American Chemical Society*, 145(49):26550–26556, 2023.

A Experimental Setups

A.1 Domain Pre-training

ChemDFM is pre-trained using the popular framework Megatron-DeepSpeed¹⁰ with Zero-2 [Rajbhandari et al., 2020] optimization technique based on LLaMa-13B [Touvron et al., 2023a]. We train ChemDFM using AdamW [Loshchilov and Hutter, 2019] with $(\beta_1, \beta_2) = (0.9, 0.95)$. During training, our model deals with 4M tokens per batch with a maximum sequence length of 6K. The maximum learning rate is 5e-5 under the cosine learning rate scheduler.

¹⁰<https://github.com/microsoft/Megatron-DeepSpeed?tab=readme-ov-file>

Table 7: Benchmark full test-set evaluation results of different models in text-based molecule design tasks. The best results among specialist and generalist models are highlighted in bold, respectively. †: reproducing results.

Model	Exact↑	BLUE↑	Dis↓	Validity↑	MACCS↑	RDK↑	Morgan↑
<i>task-specific specialist models</i>							
MolXPT [Liu et al., 2023]	21.5	-	-	98.3	0.859	0.757	0.667
Text+Chem T5 [Christofidellis et al., 2023]	32.2	0.853	16.87	94.3	0.901	0.816	0.757
Mol-Instruction [Fang et al., 2023]	0.2	0.345	41.4	100	0.412	0.231	0.147
<i>LLM-based generalist models</i>							
Galactica-30B (10-shot)†	0.3	0.295	64.3	82.2	0.356	0.239	0.186
ChemDFM-13B	43.2	0.839	16.9	97.6	0.901	0.829	0.759

B.2 Text-Based Molecule Design

B.2.1 Task Introduction

The test set of ChEBI-20 is also exploited for this task in ChemLLMBench. Models are asked to predict the SMILES of the molecule that fits the given description. Two kinds of metrics are utilized to measure the performance of this task. The first set of metrics measures the text-based similarity of the predicted SMILES compared to the golden SMILES, which includes exact match, BLUE, and Levenshtein distance. The second set of metrics measures the chemical similarity of the predicted molecules compared to the golden molecules. That is mainly composed of the validity of the predicted SMILES and the FTS (fingerprint Tanimoto Similarity) [Tanimoto, 1958] in terms of MACCS [und David Metzener, 1988], RDK¹¹, Morgan [Morgan, 1965].

B.2.2 Prompt Format

We use the same prompt introduced in Guo et al. [2023].

B.2.3 Additional Results

To achieve a fair comparison with task-specific specialist models, we evaluate the performance of ChemDFM on the full test set of ChEBI-20 on this task. The results are illustrated in Table 7. ChemDFM surpasses the performance of the advanced specialist models on the major metrics while achieving comparable performance on others. Specifically, ChemDFM outperforms the specialist models on exact match scores and all three FTS-based similarity scores, which indicates that ChemDFM can make more reliable predictions based on the descriptions compared with specialist models.

B.3 Molecular Property Prediction

B.3.1 Task Introduction

The molecular property prediction tasks in ChemLLMBench consist of five tasks from MoleculeNet benchmark [Wu et al., 2018], including BACE, BBBP, HIV, ClinTox, and Tox21. Among them, BACE and BBBP are each a balanced binary classification task. HIV is an unbalanced binary classification task. ClinTox and Tox21 comprise two and twenty-one unbalanced binary classification tasks, respectively.

B.3.2 Prompt Format

We use the same prompts introduced in Guo et al. [2023].

B.3.3 Additional Results

During evaluations, we leverage a popular and more challenging dataset split provided by DeepChem library [Ramsundar et al., 2019]. We reproduce the results of the baseline models, including GPT-4, LLaMa-2-13B-chat, and Galactica (30B). Apart from the results in the Quantitative Evaluation

¹¹<https://www.rdkit.org/>

Table 8: AUC-ROC scores [Bradley, 1997] of different models under different settings in molecular property prediction tasks. †reproducing results (The results of GPT-4 were obtained in January 2024).

Model	BACE↑	BBBP↑	ClinTox↑	HIV↑	Tox21↑
<i>LLM-based generalist models</i>					
GPT-4 (0-shot)†	62.5	61.5	51.6	65.9	55.2
GPT-4 (8-shot)†	45.9	61.8	59.3	50.8	60.6
LLaMa-2-13B-chat (0-shot)†	26.0	60.3	45.7	29.0	51.7
LLaMa-2-13B-chat (8-shot)†	72.9	52.3	42.1	70.8	45.9
Galactica-30B [Taylor et al., 2022]	72.7	59.6	82.2	75.9	68.5
ChemDFM-13B (0-shot)	<u>78.4</u>	<u>66.7</u>	89.9	<u>73.6</u>	79.8
ChemDFM-13B (8-shot)	81.7	67.9	<u>85.3</u>	73.3	<u>76.7</u>

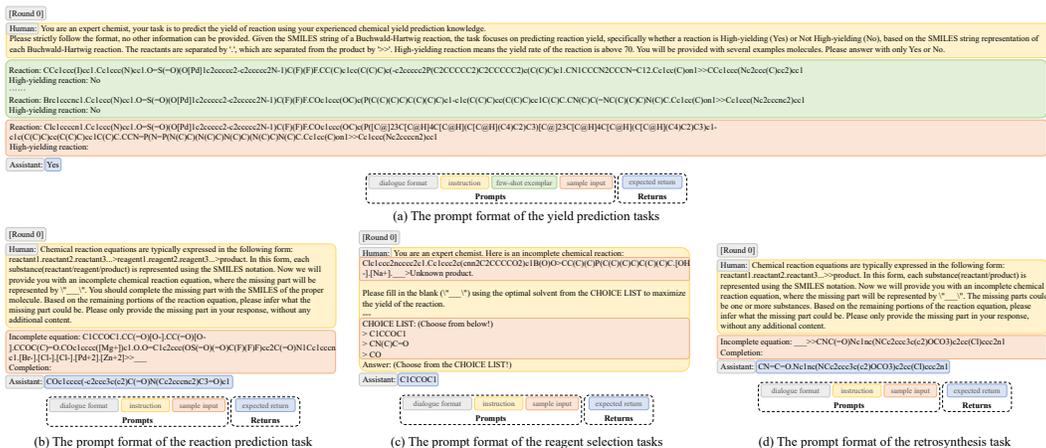


Figure 7: Prompt format of the reaction prediction and retrosynthesis tasks

Section of the main text, we also conduct few-shot experiments. The results are shown in Table 8. It is worth noticing that the performances under the few-shot settings are not always better than those under the zero-shot setting. That may be a result of the scaffold-vertical dataset split we use in our experiments. Because under the scaffold-vertical setting, the exemplars provided by the training split may be less helpful for the test samples.

B.4 Reaction Prediction and Retrosynthesis

B.4.1 Task Introduction

In ChemLLMBench, there are four types of tasks targeted at evaluating models' capabilities of reaction understanding. The yield prediction tasks ask models to predict whether the given reaction is a high-yield reaction and are constructed based on two High-Throughput experimentation (HTE) datasets: the Buchwald-Hartwig dataset [Ahneman et al., 2018] and the Suzuki-Miyaura dataset [Reizman et al., 2016]. The reaction prediction task asks the model to predict the product of the given reaction. ChemLLMBench utilizes the USPTO-MIT dataset [Jin et al., 2017] for this task. The reagent selection tasks focus on selecting the reagent that can maximize the yield of the reaction from a list of candidates. ChemLLMBench constructs three reagent selection tasks based on the dataset proposed by Perera et al. [2018]. The retrosynthesis task focuses on predicting the reactants of the given reactions and is constructed based on the USPTO-50K dataset [Schneider et al., 2016]. Accuracy is utilized to measure the performances except for the ligand selection task which uses top 50% accuracy.

Table 9: Accuracy scores of different models in yield prediction tasks. B-H and Suzuki stand for the Buchwald-Hartwig dataset and the Suzuki-Miyaura dataset, respectively. †: results from Guo et al. [2023]. *: 8-shot results

Model	B-H	Suzuki
<i>task-specific specialist models</i>		
UAGNN [Kwon et al., 2022]	96.5	95.7
<i>LLM-based generalist models</i>		
GPT-4 ^{†*}	<u>80.0</u>	<u>76.4</u>
LLaMa-2-13B-chat [†]	0.8	0.6
Galactica-30B [†]	0.0	0.8
ChemDFM-13B	82.7	79.3

Table 10: Benchmark results of different models in reaction prediction tasks. †: results from Guo et al. [2023].

Model	Accuracy	Validity
<i>task-specific specialist models</i>		
Chemformer [Irwin et al., 2022]	93.8	100
Mol-Instruction [Fang et al., 2023]	4.5	100
InstructMol [Cao et al., 2023a]	53.6	100
<i>LLM-based generalist models</i>		
GPT-4 (20-shot) [†]	<u>23.0</u>	93.0
LLaMa-2-13B-chat (20-shot) [†]	3.2	72.2
Galactica-30B (5-shot) [†]	3.6	<u>94.8</u>
ChemDFM-13B (0-shot)	49.0	98.0

B.4.2 Prompt Format

We reformat the prompt provided by Guo et al. [2023] using the SMILES notations for reactions. Specifically, the examples of our prompts are illustrated in Figure 7.

B.4.3 Additional Results

The complete results for the yield prediction tasks, the reaction prediction task, the reagent selection tasks, and the retrosynthesis tasks are shown in Table 9, Table 10, Table 12, and Table 11, respectively.

C More Qualitative Analysis

C.1 Paper Reading

We first test the models with questions that only involve known knowledge (Figure 8).

Q-S1 (Q1) is an example of knowledge-intensive questions. Models only need to memorize the details and mechanisms of Catellani-type reactions [Catellani et al., 1997] to answer the question correctly. The key point of the answer to this question is “**regioselectivity**”. While Galactica can hardly answer the question and LLaMa-2 misses the key point of the answer, ChemDFM accurately captures the key point to answer the question and provides a comprehensive answer. GPT-4 gives the best reply as it not only points out “regioselectivity” but also gives the result of the regioselectivity of norbornene. ChemDFM is the only model that tries to provide a detailed description of the mechanism behind the reaction. However, it makes minor mistakes when doing so.

Q-S2 asks for the regioselectivity of the Diels-Alder reaction [Kloetzel, 1948]. Only ChemDFM successfully answers the key points to this question, which is the result of the regioselectivity. GPT-4 provides a detailed introduction to the Diels-Alder reaction and regioselectivity but fails to answer the specific regioselectivity of the Diels-Alder reaction, while LLaMa-2 only gives the factors that could influence the regioselectivity. They do not answer the question.

LLMs	Q-S1 (Q1): Please briefly introduce the effect of norbornene in Catellani-type reactions.	Q-S2: What is the regioselectivity of Diels-Alder reaction?	Q-S3: What is Dess-Martin periodinane used for?
GP4	Norbornene plays a crucial role in Catellani-type reactions as a mediator or a template. It helps to control the ortho-allylation and palladium sites, enabling the palladium to move from one ortho position to another. This movement allows for the selective functionalization of arenes at the ortho and ipso positions. Without norbornene, the reaction would not proceed in the desired manner, leading to different products. Therefore, norbornene is essential for the success of Catellani-type reactions.	Regioselectivity in the Diels-Alder reaction refers to the preference for formation of one constitutional isomer over another. It is determined by the relative stability of the transition state leading to the different products. The reaction is typically controlled by the electronic and steric properties of the diene and dienophile. The most common outcome is the formation of a product where electron-rich areas of the diene interact with electron-poor areas of the dienophile. This is often referred to as the "normal electron demand" Diels-Alder reaction.	Dess-Martin periodinane is a chemical reagent used primarily for the oxidation of primary and secondary alcohols to aldehydes and ketones, respectively. It is popular in organic chemistry due to its high selectivity, mildness, and ease of use.
ColBERTv2.0	The norbornene is used to generate a 7 membered palladacycle. The norbornene is not a substrate in the reaction, but rather a reagent that is used to generate the catalyst. The norbornene is not consumed in the reaction, and is not a product.	The regioselectivity of the Diels-Alder reaction is determined by the stability of the diene and the dienophile. The diene is usually the more stable of the two.	Dess-Martin periodinane is used to oxidize alcohols to aldehydes or ketones. It is also used to oxidize primary alcohols to carboxylic acids.
LLMs v2.1 (B3) chat	Norbornene is a cyclic molecule that has been widely used as a catalyst in Catellani-type reactions. The presence of norbornene in these reactions has been found to have a significant effect on the reaction mechanism and the final product distribution. In general, norbornene acts as a "cage" around the reacting molecules, facilitating their interaction and promoting the formation of the desired products. The cyclic structure of norbornene also allows for the formation of complexes with the reacting molecules, which can enhance the reaction rate and selectivity.	The regioselectivity of the Diels-Alder reaction is a critical aspect of the reaction, as it determines which carbon atoms of the diene participate in the reaction and which ones remain unreacted. The regioselectivity of the Diels-Alder reaction is influenced by several factors, including the type of diene and the type of dienophile used, as well as the reaction conditions. In general, the Diels-Alder reaction is highly regioselective, meaning that it preferentially forms one or a few major products over a wide range of possible minor products. This is due to the fact that the diene and dienophile molecules have a high degree of complementarity, which allows them to form a stable and energetically favorable complex. There are several factors that influence the regioselectivity of the Diels-Alder reaction: 1. Electronic effects: The electronic properties of the diene and dienophile molecules can influence the regioselectivity of the reaction. For example, if the diene has a high degree of electron negativity, it may preferentially form bonds with the more electronegative atoms of the dienophile. 2. Steric effects: The size and shape of the diene and dienophile molecules can also influence the regioselectivity of the reaction. For example, if the diene is bulky, it may have difficulty forming bonds with the more sterically hindered atoms of the dienophile. 3. Solvent effects: The solvent in which the Diels-Alder reaction is carried out can also influence the regioselectivity of the reaction. For example, some solvents may favor the formation of one particular product over others. 4. Temperature and pressure: The reaction conditions, such as temperature and pressure, can also influence the regioselectivity of the reaction. For example, higher temperatures and pressures may favor the formation of more highly substituted products. In general, the regioselectivity of the Diels-Alder reaction is influenced by a combination of these factors, and the exact regioselectivity of the reaction can vary depending on the specific diene and dienophile molecules being used. However, the reaction is highly regioselective, and the major product is usually formed with high yield and purity.	Dess-Martin periodinane (DMP) is a highly effective and versatile reagent used in organic synthesis for the oxidation of various functional groups. It is a combination of periodinane (DCC-CH(OAc)H) and 4-dimethylaminopyridine (DMAP), which enhances the reactivity of the periodinane moiety. DMP is commonly used for the following reactions: 1. Epoxidation: DMP is a powerful oxidizing agent for the epoxidation of unsaturated compounds, such as alkenes and alkynes. It can introduce an epoxide group at the carbon-carbon double bond, which is useful for further synthetic modifications. 2. Hydroxylation: DMP can hydroxylate various substrates, including aromatic and aliphatic compounds, at the positions bearing an electron-withdrawing group. This reaction is particularly useful for the synthesis of alcohols, carboxylic acids, and amides. 3. Alkylation: DMP can alkylate various substrates, including aromatic and aliphatic compounds, at the positions bearing an electron-rich group. This reaction is useful for the synthesis of alkylated compounds, which are important intermediates in the production of pharmaceuticals and other chemicals. 4. Cyclization: DMP can induce cyclization reactions, such as the formation of lactones and lactams, by oxidizing the substrate at the carbon-carbon double bond. This reaction is useful for the synthesis of complex molecules, such as natural products and other chemicals. 5. Fluorination: DMP can also be used for the fluorination of various substrates, including aromatic and aliphatic compounds. This reaction is useful for the synthesis of fluorinated compounds, which are important in the production of pharmaceuticals and other chemicals.
ChemBERTa-1.1.0	Norbornene is a key component in Catellani-type reactions, which are a class of palladium-catalyzed cross-coupling reactions that allow for the efficient formation of carbon-carbon bonds. In these reactions, norbornene serves as a transient directing group, which helps to control the regioselectivity of the reaction and enables the formation of challenging cross-coupled products that are difficult to access using traditional cross-coupling methods. Specifically, norbornene undergoes oxidative addition to a palladium(0) catalyst to form a palladium(II) complex, which can then react with an organic halide and a boronic acid or ester to form the cross-coupled product. The norbornene group is then removed from the product under acidic conditions, regenerating the palladium(0) catalyst and completing the catalytic cycle.	The Diels-Alder reaction is a [4 + 2] cycloaddition reaction between a diene and a dienophile. The regioselectivity of this reaction depends on the structure of the diene and dienophile, as well as the substituents present. Generally, the reaction proceeds with the substituents on the diene and dienophile being located at the 1- and 4-positions of the resulting cyclohexene. However, in some cases, other regioisomers may be formed.	Dess-Martin periodinane is a reagent commonly used in organic chemistry for the oxidation of alcohols to aldehydes and ketones. It is particularly useful for the oxidation of secondary alcohols to ketones, which can be challenging to achieve with other oxidizing agents. Dess-Martin periodinane is a relatively mild reagent, which makes it useful for the oxidation of sensitive compounds. It is also used in the synthesis of complex natural products and pharmaceuticals.

Figure 8: Examples of paper reading where only widely known domain knowledge is involved. correct and relevant information in the replies is marked in green, correct but irrelevant information in yellow, and wrong information in red. Key points of the answer are marked in bold.

LLMs	Q-S4 (Q2): What is the use of NaH in the reaction between <chem>O=C1C(=O)C(=O)C(=O)C1</chem> and <chem>CC(=O)OCC(=O)O</chem> ? <i>[Wu et al., 2023]</i>	Q-S5: What is the use of DIBAL-H to react with <chem>C=C1C(=O)C(=O)C(=O)C1</chem> ? <i>[Wu et al., 2023]</i>	Q-S6: How does <chem>C1C(=O)C(=O)C(=O)C1</chem> react with <chem>COc1ccc(NC(=O)C(=O)C(=O)C1)cc1</chem> ? <i>[Wang et al., 2022]</i>
GP4	NaH, or Sodium Hydride, is a strong base that is often used in organic chemistry as a deprotonating agent. In the reaction between <chem>O=C1C(=O)C(=O)C(=O)C1</chem> and <chem>CC(=O)OCC(=O)O</chem> , NaH would likely be used to deprotonate one of the reactants, making it more nucleophilic and therefore more reactive. This would allow it to react more readily with the other reactant, facilitating the overall reaction.	DIBAL-H (Diisobutylaluminum hydride) is a reducing agent commonly used in organic chemistry. It is particularly useful for the reduction of esters, carboxylic acids, and amides to aldehydes. In the case of the compound you provided, DIBAL-H would likely be used to reduce the ester group (COOEt) to an aldehyde. This would result in the formation of an aldehyde group (-CHO) and an alcohol group (-OH) in place of the ester. However, the structure you provided seems to be a bit complex and may not be entirely accurate. It would be helpful to have a more accurate structure to provide a more precise answer.	This reaction is a complex organic synthesis involving multiple steps. Here is a simplified version of the reaction: 1. The first step is a nucleophilic attack by the nitrogen atom of the amine group on the carbonyl carbon of the second molecule. This forms a tetrahedral intermediate. 2. The intermediate then undergoes a rearrangement, leading to the formation of a new carbon-nitrogen bond and the release of a water molecule. 3. The newly formed amide undergoes a cyclization reaction, forming a new ring structure. 4. The final step is a dehydration reaction, where a water molecule is removed, leading to the formation of the final product. Please note that this is a simplified version of the reaction and the actual mechanism may involve additional steps and intermediates. The stereochemistry of the reaction is also not detailed in this explanation.
ColBERTv2.0	The use of <chem>NaH</chem> is to deprotonate the carboxylic acid. The resulting carboxylate is a better leaving group than the carboxylic acid.	The use of DIBAL-H is to reduce the double bond.	The first step is the formation of a <chem>[carbene]</chem> https://en.wikipedia.org/wiki/Carbene . The carbene is then attacked by the oxygen of carbon monoxide. The resulting <chem>[acyl]</chem> https://en.wikipedia.org/wiki/Acyl_group is then attacked by the nitrogen of the amide.
LLMs v2.1 (B3) chat	NaH is used as a reducing agent in the reaction between <chem>O=C1C(=O)C(=O)C(=O)C1</chem> and <chem>CC(=O)OCC(=O)O</chem> . NaH would likely be used to deprotonate one of the reactants, making it more nucleophilic and therefore more reactive. This would allow it to react more readily with the other reactant, facilitating the overall reaction.	DIBAL-H (diisobutylaluminum hydride) is a strong reducing agent that is commonly used to react with carbon-carbon double bonds (C=C) and carbon-carbon triple bonds (C≡C) in organic compounds. In the case of the molecule you provided, <chem>C=C1C(=O)C(=O)C(=O)C1</chem> , DIBAL-H can be used to react with the carbon-carbon triple bond (C≡C) and the carbon-carbon double bond (C=C) present in the molecule. The reaction of DIBAL-H with the C≡C triple bond in the molecule would proceed through the formation of a metal complex, where the aluminum center of DIBAL-H coordinates with the carbon atoms of the double bond. This would result in the reduction of the triple bond to a double bond, leading to the formation of a new C=C double bond. Similarly, the reaction of DIBAL-H with the C=C double bond in the molecule would also proceed through the formation of a metal complex, where the aluminum center of DIBAL-H coordinates with the carbon atoms of the double bond. This would result in the reduction of the double bond to a single bond, leading to the formation of a new C-C single bond.	The reaction of <chem>C1C(=O)C(=O)C(=O)C1</chem> with <chem>COc1ccc(NC(=O)C(=O)C(=O)C1)cc1</chem> involves multiple steps and reactions. Here's a simplified outline of the major steps involved in the reaction: Step 1: Formation of a carbon-carbon bond The first step in the reaction is the formation of a carbon-carbon bond between the carbon atom of <chem>C1C(=O)C(=O)C(=O)C1</chem> and the carbon atom of <chem>COc1ccc(NC(=O)C(=O)C(=O)C1)cc1</chem> . This forms a new carbon-carbon bond, which is shown in the figure as a double-headed arrow. Step 2: Formation of a nitrogen-carbon bond The next step is the formation of a nitrogen-carbon bond between the nitrogen atom of <chem>C1C(=O)C(=O)C(=O)C1</chem> and the carbon atom of the newly formed carbon-carbon bond. This forms a new nitrogen-carbon bond, which is shown in the figure as a single-headed arrow. Step 3: Formation of a carbon-oxygen bond The final step is the formation of a carbon-oxygen bond between the carbon atom of the newly formed carbon-carbon bond and the oxygen atom of <chem>COc1ccc(NC(=O)C(=O)C(=O)C1)cc1</chem> . This forms a new carbon-oxygen bond, which is shown in the figure as a single-headed arrow. Step 4: Formation of a nitrogen-oxygen bond The final step is the formation of a nitrogen-oxygen bond between the nitrogen atom of <chem>C1C(=O)C(=O)C(=O)C1</chem> and the oxygen atom of the newly formed carbon-oxygen bond. This forms a new nitrogen-oxygen bond, which is shown in the figure as a single-headed arrow. Overall, the reaction of <chem>C1C(=O)C(=O)C(=O)C1</chem> with <chem>COc1ccc(NC(=O)C(=O)C(=O)C1)cc1</chem> involves the formation of a new carbon-carbon bond, a new nitrogen-carbon bond, a new carbon-oxygen bond, and a new nitrogen-oxygen bond. The resulting molecule is <chem>COc1ccc(NC(=O)C(=O)C(=O)C(=O)C1)cc1</chem> .
ChemBERTa-1.1.0	The use of NaH in the reaction between <chem>O=C1C(=O)C(=O)C(=O)C1</chem> and <chem>CC(=O)OCC(=O)O</chem> is to deprotonate the acidic alpha-hydrogen of the keto group at <chem>CC(=O)OCC(=O)O</chem> , resulting in the formation of an enolate anion. This anion can then react with the end group of <chem>O=C1C(=O)C(=O)C(=O)C1</chem> to form a new C-C bond and produce the product.	DIBAL-H is a reducing agent that can be used to reduce carbonyl groups to the corresponding alcohols. In the case of 2-(2-methoxyethoxy)-2-methyl-1-oxo-10-azapropyl-5-thioxo-4-vinylidene-1,1,1,3,3-tetraethylcyclohexane-6-carboximide, DIBAL-H can be used to reduce the carbonyl group of the amide to the corresponding aldehyde. This reduction is a key step in the synthesis of the anti-cancer drug, Piflutin.	The reaction involves the formation of a covalent bond between the ketone group of the benzimidazole aromatic compound and the amino group of the (2S)-1-(2R)-2-phenyl-2-propyl-3,4-dihydro-1H-isooquinoline-8-carboxamide. The stereochemistry of the resulting product is consistent with the reaction of the (R)-enantiomer of the benzimidazole aromatic compound with the (2S)-enantiomer of the (2R)-1-(2R)-2-phenyl-2-propyl-3,4-dihydro-1H-isooquinoline-8-carboxamide.

Figure 9: Examples of paper reading where new molecules and reactions are involved. correct and relevant information in the replies is marked in green, correct but irrelevant information in yellow, and wrong information in red. Key points of the answer are marked in bold.

Table 11: Benchmark results of different models in retrosynthesis tasks. †: results from Guo et al. [2023].

Model	Accuracy	Validity
<i>task-specific specialist models</i>		
Chemformer [Irwin et al., 2022]	53.6	100
<i>LLM-based generalist models</i>		
GPT-4 (5-shot) [†]	<u>11.4</u>	89.0
LLaMa-2-13B-chat (20-shot) [†]	0.0	72.8
Galactica-30B (5-shot) [†]	1.6	94.8
ChemDFM-13B (0-shot)	12.0	<u>91.0</u>

Table 12: Benchmark results of different models in reagent selection tasks. We report the result in accuracy scores except for Ligand Selection where we report the top 50% accuracy score. †: results from Guo et al. [2023].

Model	Reactant	Solvent	Ligand
<i>LLM-based generalist models</i>			
GPT-4 [†]	29.9	52.6	53.4
LLaMa-2-13B-chat [†]	14.5	5.0	28.4
Galactica-30B [†]	10.7	10.4	3.0
ChemDFM-13B	<u>24.0</u>	<u>12.0</u>	<u>35.0</u>

As for **Q-S3**, ChemDFM, Galactica, and GPT-4, all capture the key point to the answer (“**the oxidation of alcohols to aldehydes and ketones**”), while ChemDFM and GPT-4 further answer more properties of the Dess-Martin periodinane [Dess and Martin, 1983]. LLaMa-2, on the other hand, gives numerous wrong arguments and misses the key points.

Then, we ask the models about new molecules and new reactions which are published after January 2022. In this way, we can ensure minimal risk of data leakage and evaluate the models’ capability to handle unforeseen situations. The results are shown in Figure 9 and Figure 10.

Q-S4 (Q2) is constructed based on Yin et al. [2023]. Because the reaction mentioned in the question is a novel instance, models need to correctly identify the reaction and discover the mechanisms of it before answering the question. In practice, Galactica successfully identifies the key point of the answer, “**deprotonate**”, but fails to provide other useful information. LLaMa-2, in its reply, fails to identify the reaction mentioned in the question. Most of the information about NaH in its reply is correct but irrelevant to the reaction. GPT-4 identifies the key point of the answer but only gives a rough description of the mechanism of how it works. ChemDFM not only correctly identifies the key point of the answer but also provides an almost correct description of the mechanism.

Q-S5 is also constructed based on Yin et al. [2023]. All the models can recognize the DIBAL-H as a reducing agent, which is existing knowledge. However, only ChemDFM successfully identifies the reaction site of the new molecule, indicating its strong capabilities to handle unforeseen situations where new molecules and reactions are involved. The main mistake that ChemDFM makes is providing the wrong IUPAC name, which is a challenging task for LLMs even as a separate task (see Table 2 in the main text).

Q-S6 is constructed based on Wang et al. [2023a] and asks directly for the mechanism of the given reaction. Among the answers, the answer of ChemDFM is the most precise. Galactica and LLaMa-2 give nearly no correct information. Although GPT-4’s answer contains the correct reaction process, it also contains auxiliary processes that do not happen during the reaction, which masks the whole mechanism predicted by GPT-4 wrong. ChemDFM answers the correct reaction process with no excess. The only mistakes ChemDFM makes are again providing the wrong IUPAC names, which is a challenging task for LLMs even as a separate task (see Table 2 in the main text).

We also ask several questions focusing more on molecules and less on reactions.

LLM	<p>Q-S7 (Q3): I want to enhance the acidity of catalyst</p> <p><chem>C=C([@H]N[C@@H]2CC[C@@H]3[C@@H]([C@@H]2)N[C@@H]3)C(=O)C(F)(F)F</chem> How should I modify the catalyst? (Dargó et al., 2023)</p>	<p>Q-S8: I want to change the chiral environment of catalyst</p> <p><chem>CN1CCCC(N2C=CC=CC=C2)C1</chem> How should I modify the catalyst? (Hao et al., 2023)</p>	<p>Q-S9: Which part of</p> <p><chem>O=C1C=CC(=O)C=C(C)C=C1</chem> has the chance to connect with palladium? (Wang et al., 2023a)</p>
GPT-4	<p>To enhance the acidity of a catalyst, you can introduce more electronegative atoms or functional groups that can donate protons. In your case, you can consider the following modifications:</p> <ol style="list-style-type: none"> Introduce more sulfonic acid groups (-SO₃H): Sulfonic acid groups are highly acidic and can significantly increase the acidity of your catalyst. You can replace some of the existing functional groups with -SO₃H. Add more halogen atoms: Halogens are highly electronegative and can increase the acidity of the catalyst. You already have fluorine atoms in your catalyst. You can consider adding more fluorine atoms or other halogens like chlorine or bromine. Introduce carboxylic acid groups (-COOH): Carboxylic acid groups are also highly acidic. You can replace some of the existing functional groups with -COOH. Introduce phosphoric acid groups (-PO₃H₂): Nitro groups are also highly electronegative and can increase the acidity of the catalyst. <p>Remember, the modification should not affect the catalyst's ability to catalyze the reaction. It's always a good idea to consult with a chemist or a chemical engineer before making any modifications to a catalyst.</p>	<p>The chiral environment of a catalyst can be modified by changing the chiral center or introducing a new chiral center. The chiral center is the atom in a molecule that is attached to four different groups, usually carbon.</p> <p>In the given catalyst, the chiral center seems to be the carbon atom in the cyclohexane ring (C1(CCC)). To modify the chiral environment, you can:</p> <ol style="list-style-type: none"> Change the groups attached to the chiral center: You can replace one or more of the groups attached to the chiral center with different groups. This will result in a different arrangement of groups around the chiral center, thus changing the chiral environment. Introduce a new chiral center: You can add a new chiral center to the molecule by attaching a carbon atom to four different groups. This will create a new chiral environment in the molecule. Use a different chiral catalyst: If modifying the existing catalyst is not feasible, you can consider using a different chiral catalyst that provides the desired chiral environment. <p>Remember, any changes to the catalyst should not affect its ability to catalyze the reaction. It's always recommended to consult with a chemist or a chemical engineer before making any changes to a catalyst.</p>	<p>The part of the molecule that has the chance to connect with palladium is the C=C double bond in the (C=C) section. Palladium is often used as a catalyst in reactions involving double bonds, such as in the Heck reaction or Suzuki coupling.</p>
Galactica-V0B	<p>The catalyst you have is a Bronsted acid (https://en.wikipedia.org/wiki/Bronsted_acid). The acidity of the catalyst is determined by the strength of the conjugate base. The stronger the conjugate base, the stronger the acid. The conjugate base of your catalyst is the carbocation on the nitrogen. The stronger the carbocation, the stronger the acid. The carbocation is stabilized by the adjacent double bond. The more electron-withdrawing the double bond, the stronger the acid. The double bond is conjugated with the nitrogen lone pair. The more electron-withdrawing the nitrogen lone pair, the stronger the acid. The nitrogen lone pair is stabilized by the adjacent double bond.</p> <p>The catalyst you provided is a complex molecule with multiple functional groups, and enhancing its acidity may require careful modification. Here are some strategies that may help:</p> <ol style="list-style-type: none"> Introduction of acidic functional groups: You can introduce acidic functional groups such as carboxylic acids, sulfonic acids, or phosphoric acids into the catalyst molecule. These groups can increase the acidity of the catalyst and enhance its ability to activate the substrate. Substitution with acidic groups: You can substitute some of the non-acidic functional groups in the catalyst with acidic groups. For example, you can replace some of the hydrogen atoms with acidic groups such as hydroxyl (-OH), carboxyl (-COOH), or sulfonic acid (-SO₃H) groups. Increase the number of acidic sites: You can increase the number of acidic sites in the catalyst by introducing more acidic functional groups or by increasing the number of existing acidic sites. This can be done by modifying the molecular structure of the catalyst to increase the number of carboxylic acid, sulfonic acid, or phosphoric acid groups. Use of acidic additives: You can add acidic additives such as sulfuric acid, hydrochloric acid, or acetic acid to the reaction mixture to enhance the acidity of the catalyst. These additives can increase the concentration of acidic sites in the reaction mixture and improve the catalytic activity of the catalyst. Modification of the support: The support material can also affect the acidity of the catalyst. You can modify the support material to increase its acidity. For example, you can use a support material that is more acidic, such as silica gel or alumina, or you can modify the support material with acidic functional groups such as carboxylic acids or sulfonic acids. <p>It is important to note that modifying the catalyst to enhance its acidity can also affect its other properties, such as its stability and selectivity. Therefore, it is important to carefully optimize the modification of the catalyst to achieve the desired balance of properties. Additionally, the optimal modification of the catalyst may depend on the specific reaction conditions and the substrate being used, so it may be necessary to perform experiments under different conditions to find the best modification strategy.</p>	<p>The catalyst you have is a diaminomethylbenzimidazolium cation (https://en.wikipedia.org/wiki/Diaminomethylbenzimidazolium_cation). The chiral environment is created by the two amino groups. You can change the chiral environment by changing the amino groups.</p>	<p>The palladium catalyst is used to form the carbon-carbon bond between the two molecules. The carbon-carbon bond is formed between the carbon in the alkyne and the carbon in the arene. The carbon in the alkyne is the one that is attached to the nitrogen.</p>
LLaMa-2-70B	<p>To change the chiral environment of the catalyst, you could consider modifying the chiral substituents on the nitrogen atoms.</p> <p>One approach could be to introduce a chiral auxiliary group, such as a quaternary amine or a chiral sulfonamide, on one of the nitrogen atoms. This would allow you to control the stereochemistry of the reaction and potentially improve the enantioselectivity of the catalyst.</p> <p>Another approach could be to modify the ligand environment around the metal center. For example, you could introduce a chiral ligand, such as a chiral amine or a chiral phosphine, that coordinates to the metal center and influences the stereochemistry of the reaction.</p> <p>It is also possible to modify the metal center itself to introduce chirality. For example, you could use a chiral metal such as a chiral transition metal or a chiral main group element, or you could introduce a chiral ligand that coordinates to the metal center.</p> <p>It is important to note that modifying the catalyst can potentially affect its activity and selectivity, so it may be necessary to optimize the modified catalyst to achieve the desired properties. Additionally, the chiral environment of the catalyst may be sensitive to the conditions of the reaction, such as the solvent, temperature, and reaction time, so it may be necessary to optimize these conditions as well.</p> <p>Overall, modifying the chiral environment of the catalyst can be a powerful way to improve the enantioselectivity of the reaction, but it requires careful consideration of the specific modifications and their potential effects on the catalyst's activity and selectivity.</p>	<p>The molecule you provided, <chem>O=C1C=CC(=O)C=C(C)C=C1</chem>, contains several potential binding sites for palladium. Here are some possible ways that palladium could connect with this molecule:</p> <ol style="list-style-type: none"> Binding to the nitrogen atom: The nitrogen atom in the molecule is electronegative and could potentially form a coordinate covalent bond with a palladium center. Binding to the carbon atoms: The carbon atoms in the molecule are sp³ hybridized and could potentially form a coordinate covalent bond with a palladium center. Binding to the oxygen atoms: The oxygen atoms in the molecule are electronegative and could potentially form a coordinate covalent bond with a palladium center. Binding to the sulfur atom: The sulfur atom in the molecule is electronegative and could potentially form a coordinate covalent bond with a palladium center. <p>It's difficult to predict the exact binding site for palladium on this molecule without knowing the specific conditions of the reaction, but these potential binding sites could be important for the catalytic activity of the palladium center.</p> <p>It's worth noting that the molecule you provided is a complex molecule with many functional groups and atoms that could potentially participate in the catalytic process. The exact binding site for palladium may depend on the specific reaction conditions and the conformation of the molecule.</p>	
ChemDFM	<p>To enhance the acidity of the catalyst, you could consider introducing electron-withdrawing groups on the aromatic ring of the <chem>[C@@H]1N[C@@H](C(F)(F)F)C(F)(F)F1</chem> moiety. This could draw electron density away from the nitro group and increase its acidity. Some examples of such groups could be fluoro, chloro, or bromo substituents. However, it is important to note that the introduction of these groups could also affect the binding affinity of the catalyst to its target, so careful optimization would be required.</p>	<p>To change the chiral environment of the catalyst, you will need to modify the structure of the N-aryl substituent. This can be done by introducing a chiral element into the aryl group, such as a chiral center or a chiral substituent. You may also need to adjust the connectivity of the aryl group to the imidazole moiety to achieve the desired chirality. It is important to note that changing the chiral environment of the catalyst may also affect its activity and selectivity, so you may need to optimize the new structure to achieve the desired results.</p>	<p>The part of the molecule that has the greatest chance of connecting with palladium is the amide bond between the amine moiety and the rest of the molecule. This bond is formed between the carbonyl carbon of the amine and the nitrogen atom of the rest of the molecule, and is a common site for coordination with palladium.</p>

Figure 10: Examples of paper reading where new molecules and reactions are involved. **correct** and **relevant information** in the replies is marked in green, **correct but irrelevant information** in yellow, and **wrong information** in red. **Key points of the answer** are marked in bold.

Q-S7 (Q3), constructed based on Dargó et al. [2023], focus on the modification of catalyst molecules. The molecule mentioned in the question is a novel instance and models need to infer the chemical properties of that molecule to answer the question. The key point of the answer is “**introducing electron-withdrawing groups on the aromatic rings**” as this method has the potential to increase the acidity while keeping the catalytic ability of the molecule. Among the LLMs, only ChemDFM successfully answers the key point, while others either fail to provide any specific solutions or give wrong solutions which will damage the catalytic ability of the molecule.

Q-S8, constructed based on Hao et al. [2023], focus on the modification of chiral environment. In the given molecule, there are two chiral centers. GPT-4 identifies the point chiral center and provides candidate methods that are not all correct. The other three models identify the axial chirality which is challenging to identify with only the SMILES notation. Among the three models, Galactica gives almost no detailed method to change the chiral environment, LLaMa-2 gives one correct method with more wrong ones, and ChemDFM provides two correct methods one of which is targeting specifically the axial chirality.

Q-S9 is constructed on Wang et al. [2023a]. It asks for the coordinate sites between the given molecule and palladium. There are a total of three coordinate sites. GPT-4 and ChemDFM each identify one of them, while Galactica and LLaMa-2 fail to identify any.

C.2 Dialogue-Based Human-AI Collaboration

We demonstrate two more examples of dialogue-based human-AI collaboration based on ChemDFM here. The dialogues are also inspired by the recently published papers to minimize the risk of data leakage and evaluate ChemDFM’s capability to handle unforeseen situations during human-AI collaboration.

The dialogue shown in Figure 11 is inspired by Yin et al. [2023]. During the dialogue, the human researcher first asks for the role of LiCl in the given reaction. ChemDFM successfully identifies

Round 1



What is the **role of LiCl** in the process of generating O=C1C=C2C3(CCC2)CC4(C)CCC(C)(C)C4C31C from O=C1C(C(OC)=O)=C2C3(CCC2)CC4(C)CCC(C)(C)C4C31C?

The role of LiCl in this process is to act **as a catalyst** for the reaction. LiCl is added to a solution of the reactants in dry acetone, and it helps to **facilitate the Diels-Alder reaction** between the diene and dienophile moieties of the two starting materials. In this reaction, LiCl **serves as a Lewis acid, coordinating with the carbonyl oxygen** of one of the starting materials and increasing the electrophilicity of the dienophile. This makes it more reactive towards the diene, which is then able to add to the dienophile to form the final product



Round 2



But it is **not Diels-Alder reaction**. The ester group on reactant leaves to form product.

You are correct, I apologize for the mistake in my previous response. The process you described **is not a Diels-Alder reaction, but rather a decarbonylation reaction**. In this reaction, the LiCl catalyst helps to **facilitate the loss of the carbonyl group** from one of the starting materials, **allowing the decarbonylated product to be formed**. The exact mechanism by which LiCl promotes this reaction is not entirely clear, but it is believed to **involve the formation of a carbanion intermediate that can undergo beta-elimination to remove the carbonyl group**.



Round 3



After using LiCl, a **lithium salt intermediate product** was obtained. What should be added next to obtain the final product?

After using LiCl to promote the decarbonylation reaction, the resulting intermediate product will contain a lithium salt. To obtain the final product, **the lithium salt can be hydrolyzed by adding water to the reaction mixture**. This will cause the lithium salt to precipitate out as a solid, and the resulting supernatant can be isolated and purified to obtain the final product. Alternatively, the lithium salt can be converted back into one of the starting materials by treating it with a proton source, such as acid.



Figure 11: Example showing ChemDFM as an assistant researcher in the design of experiment through free-form dialogue. **Key points of the answer** are marked in bold and **wrong information** in red. The dialogue is inspired by Yin et al. [2023].

Round 1



What **intramolecular reactions will occur** when AIBN and Bu₃SnH is added to CC(C)(C)[Si](C)(C)OC1(c2cc3c(cc2Br)OCO3)C(=O)N2CCCC23C=C(OC(=O)c2ccccc2)C(=O)C13?

When AIBN and Bu₃SnH is added to CC(C)(C)[Si](C)(C)OC1(c2cc3c(cc2Br)OCO3)C(=O)N2CCCC23C=C(OC(=O)c2ccccc2)C(=O)C13, an intramolecular reaction will occur to **form a new ring**. The reaction involves the **formation of a new C-C bond between the carbon adjacent to the carbonyl group and the carbon adjacent to the silicon group in the starting material**. This reaction is called the "ene" reaction.



Round 2



But this is a radical debromination reaction.

You are correct, I apologize for the error in my previous response. The reaction between AIBN and Bu₃SnH in the presence of CC(C)(C)[Si](C)(C)OC1(c2cc3c(cc2Br)OCO3)C(=O)N2CCCC23C=C(OC(=O)c2ccccc2)C(=O)C13 will **result in a debromination reaction**, not an ene reaction. **The debromination reaction will occur at the beta-position of the enone moiety, resulting in the formation of a new C-C bond and the release of bromine**.



Round 3



Subsequently, TBAF is added, and **what is its purpose?**

TBAF (tetrabutylammonium fluoride) is a source of fluoride ion and will be added to the reaction mixture after the debromination reaction. The purpose of adding TBAF is to **hydrolyze the silicon-containing group in the product, thereby removing the protecting group and revealing the alcohol functional group**. This step is necessary to allow for further functionalization of the molecule.



Figure 12: Example showing ChemDFM as an assistant researcher in the design of experiment through free-form dialogue. **Key points of the answer** are marked in bold and **wrong information** in red. The dialogue is inspired by Zhuang et al. [2023].

the LiCl as a catalyst while misjudging the type of the reaction. To correct the answer, the human researcher points out the key error in the answer with some important details of the reaction (which can be easily discovered by comparing the product with the reactant). ChemDFM then corrects its mistake with even more details about the reaction process. To further validate whether ChemDFM fully understands the unforeseen reaction, we continue to ask about the post-processing procedure which is necessary to get the final product. ChemDFM successfully captures the key point to the question and gives the correct answer.

The dialogue shown in Figure 12 is inspired by Zhuang et al. [2023]. ChemDFM first gives a partially correct answer to the question from the human researcher where it misjudges the position of the newly formed C-C bond and the type of the reaction. With the help of human correction, ChemDFM then realizes the mistakes and corrects them. Then the human researcher further asks about the next reaction that is conducted in Zhuang et al. [2023] without clarifying the current molecule composition of the system or restating the previous reaction. ChemDFM can infer this information from the dialogue history and correctly answer the question.

In these dialogues, ChemDFM shows promising capabilities in handling unforeseen situations, error correction, and inferring information from dialogue history. These capabilities can be attributed to the fact that ChemDFM comprehends both natural language and chemical language. This allows a universal language protocol established between ChemDFM and human researchers, enabling meaningful human-AI collaborations.