

# TinyChemVL: Advancing Chemical Vision-Language Models via Efficient Visual Token Reduction and Complex Reaction Tasks

Xuanle Zhao<sup>1,2</sup>, Shuxin Zeng<sup>1,2</sup>, Xinyuan Cai<sup>1\*</sup>, Xiang Cheng<sup>1</sup>,  
Duzhen Zhang<sup>1</sup>, Xiuyi Chen<sup>1</sup>, Bo Xu<sup>1,2\*</sup>

<sup>1</sup> The Key Laboratory of Cognition and Decision Intelligence for Complex Systems,  
Institute of Automation, Chinese Academy of Sciences

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

## Abstract

While Vision Language Models (VLMs) have demonstrated remarkable capabilities in general visual understanding, their application in the chemical domain has been limited, with previous works predominantly focusing on text and thus overlooking critical visual information, such as molecular structures. Current approaches that directly adopt standard VLMs for chemical tasks suffer from two primary issues: (i) computational inefficiency of processing entire chemical images with non-informative backgrounds. (ii) a narrow scope on molecular-level tasks that restricts progress in chemical reasoning. In this work, we propose **TinyChemVL**, an efficient and powerful chemical VLM that leverages visual token reduction and reaction-level tasks to improve model efficiency and reasoning capacity. Also, we propose **ChemRxn-V**, a reaction-level benchmark for assessing vision-based reaction recognition and prediction tasks. Directly predicting reaction products from molecular images poses a non-trivial challenge, as it requires models to integrate both recognition and reasoning capacities. Our results demonstrate that, with only 4B parameters, TinyChemVL achieves superior performance on both molecular and reaction tasks, while also demonstrating faster inference and training speeds compared to existing models. Notably, TinyChemVL outperforms ChemVLM while utilizing only 1/16th of the visual tokens. This work builds efficient yet powerful VLMs for chemical domains by co-designing model architecture and task complexity.

**Code** — <https://github.com/xxlllz/TinyChemVL>

## Introduction

The remarkable success of Large Language Models (LLMs) on diverse tasks, from question answering to code generation (Zhang et al. 2024d; Grattafiori et al. 2024), has spurred significant interest in leveraging their reasoning capabilities to accelerate scientific discovery. This trend is particularly prominent in domains like mathematics and chemistry, where core entities such as equations and molecules can be parsed in structured textual formats (Lu et al. 2022; Lightman et al. 2023; Zhang et al. 2024a, 2025). Especially in the chemical field, the molecules are parsed into SMILES (Lightman et al. 2023) and SELFIES (Krenn et al. 2022)

\*Corresponding Authors.

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

formats for the following analysis. However, representing the chemical molecules in text format inevitably omits some spatial information. More recently, Vision Language Models (VLMs) have demonstrated remarkable capabilities in addressing a wide range of visual tasks, such as image captioning and question answering (Liu et al. 2023; Zhao et al. 2025b), which can handle more complex tasks. However, in the chemical field, there is still a lack of exploration on how to efficiently utilise VLMs for chemical tasks. Recent works have demonstrated initial attempts in adapting VLMs for chemical applications such as molecule caption and property prediction (Li et al. 2025). However, the models are generally fine-tuned on general domain VLMs, which limits their efficiency in handling chemical problems.

Although previous works have made some initial explorations in the field of vision-based chemistry, several challenges remain to be addressed. (i) **Visual redundancy in processing molecule image.** In the molecular images, all structural information is contained within the graphical representations of the molecules, while the background region is considered non-informative. Consequently, there is significant redundancy in the visual tokens used to represent the entire image, as most tokens correspond to this uninformative space. Also, the partitioning of the molecular image into patches risks the loss of important structural information. For example, in the ChemVLM (Li et al. 2025), a  $800 \times 800$  image needs 1280 visual tokens for embedding, which is around 100 times the number of tokens in the textual questions, which significantly increases the time and computational cost for both training and inference. (ii) **Limited model performance and evaluation diversity.** Currently, the capacities of chemical VLMs are limited, and evaluations are not comprehensive. Existing works (Li et al. 2025; Tan et al. 2025) focus on tasks at the molecular level, such as SMILES Optical Character Recognition (OCR) and property prediction, while ignoring issues at the reaction level. Unlike direct recognition, reaction-level tasks, such as reaction prediction, require the models to perform chemical understanding and reasoning to greater extents. (iii) **Bottleneck on molecular image generation.** Recently, ChemM-LLM (Tan et al. 2025) proposes that generating the molecular image would expand the applications of VLMs in chemistry and combine VQ-GAN (Esser, Rombach, and Ommer 2021) with LLMs to generate images. However, purely gen-

erating molecular images requires external tools to parse images back into SMILES strings for subsequent analysis. Also, there is a mismatch between the architecture of current VLMs and the requirements of direct image generation.

In this work, we propose TinyChemVL, a small-scale VLM with only 4B parameters, achieving the state-of-the-art (SOTA) performance across various vision-based chemical tasks. (i) To reduce visual token redundancy, we adopt an adaptive token merge and pruning approach (Bolya et al. 2022; Zeng and Yu 2024). Utilizing the same InternVL architecture, our model reduces visual tokens to 1/16 of those in ChemVLM, significantly decreasing redundancy and capturing complex molecular structure. (ii) To advance the field of vision-based chemistry, we introduce a new benchmark named VisRxnBench. It is designed to evaluate vision-based reaction recognition and prediction tasks, providing 5,000 samples for each. To the best of our knowledge, we are the first to define the task of predicting products directly from the reactant image, which expands the scope of vision-based chemistry. Furthermore, we conduct a large-scale and diverse dataset for training. Benefiting from the reduced visual tokens, training time on large-scale data increases negligibly. (iii) To enable TinyChemVL to generate the molecular images. We employ executable code as the intermediary content for generating molecular images (Zhao et al. 2025a). This approach streamlines the process because the SMILES string is directly included within the generated code, enabling the generation molecular images and eliminating the need for additional tools in subsequent analyses. We evaluate TinyChemVL on various benchmarks, and the results show that it achieves the SOTA performance on most tasks with only 4 billion parameters. Beyond its strong performance, TinyChemVL exhibits significantly reduced training and inference times compared to existing models. In summary, the main contributions of this work are as follows:

- We introduce TinyChemVL, an efficient chemical VLM that achieves SOTA performance in vision-based chemical tasks, with only 4B parameters.
- We adopt a token merging and pruning strategy to reduce visual tokens from large background regions and merge complex structures in molecular images. This approach significantly reduces the number of visual tokens, leading to substantially faster inference and training speeds.
- We introduce VisRxnBench, a novel benchmark alongside a large-scale training dataset for evaluating vision-based, reaction-level tasks. Our evaluation shows existing VLMs perform poorly on reaction-level tasks, especially when directly predicting products from images.

## Related Works

### Chemical Language Models

Leveraging language models to solve chemical-related tasks has garnered significant attention from the research community. Previous works (Edwards, Zhai, and Ji 2021; Edwards et al. 2022) utilize pretrained SciBERT (Beltagy, Lo, and Cohan 2019) and T5 (Raffel et al. 2020) as the backbones and fine-tune on molecule-to-natural language

translation tasks. However, these works are generally task-specific. Building on the powerful generalization capabilities of LLMs, chemical LLMs demonstrate improved generalization across a wider range of chemical tasks. Recent works such as ChemLLM (Zhang et al. 2024a) and ChemDFM (Zhao et al. 2024b) achieve this by conducting large-scale chemical instruction tuning datasets and fine-tuning on existing LLMs such as InternLM. UniMoT (Zhang et al. 2024c) utilize a molecular tokenizer instead of a projection layer to encode the molecular graph information. ChemCrow (M. Bran et al. 2024) integrates expert-designed tools with GPT-4 as the backbone for chemistry tasks.

### Vision Language Models

Recently, leveraging the projection layer to align LLMs and vision encoders, VLMs have achieved superior performance on many visual tasks (Zhang et al. 2024b; Bi et al. 2024). Current models, like GPT-4o (OpenAI 2024) and InternVL (Chen et al. 2024) have shown strong performance on vision understanding and reasoning tasks. Recently, VLMs have been increasingly applied to vision-based tasks in the field of chemistry. For instance, ChemVLM (Li et al. 2025) introduces vision-based chemical datasets and utilizes InternVL architecture (Chen et al. 2024) as the backbone for fine-tuning. ChemDFM-X (Zhao et al. 2024a) utilizes distinct encoders to align representations from visual and graph modalities with the texts. In the generative domain, ChemM-LLM (Zhang et al. 2024a) employs the VQ-GAN (Esser, Rombach, and Ommer 2021) as the encoder and decoder for molecular image understanding and generation. However, these models are typically directly fine-tuned on existing architectures and suffer from significant visual token redundancy, resulting in computational inefficiency. To address these issues, we propose TinyChemVL, which employs an adaptive token merging and pruning strategy, along with a large-scale and diverse training corpus, to enhance both model efficiency and performance.

## Methods

The overall structure of TinyChemVL is illustrated in Figure 1. TinyChemVL adopts the well-established ViT-MLP-LLM architecture, following the same framework conducted by InternVL (Chen et al. 2024) and ChemVLM (Li et al. 2025) to ensure fair comparisons. In contrast to prior works, we adopt InternVL2.5-4B as our backbone. To process high-resolution images, the dynamic resolution strategy segments the image into several  $448 \times 448$  tiles based on its dimensions and concatenates them with the downsampled image.

### Visual Token Redundancy

The vision transformer aims to encode the image into visual features. Given the input image resolution  $N \times N$ , the ViT encode the image to  $(N//P)^2$  tokens with  $P \times P$  patch size. When facing the high-resolution image, the dynamic resolution strategy first segments it into standard-resolution tiles, followed by a directly down-sampled thumbnail. Then the ViT encode each tile independently and concatenates the processed tokens to conduct the visual tokens. We identify

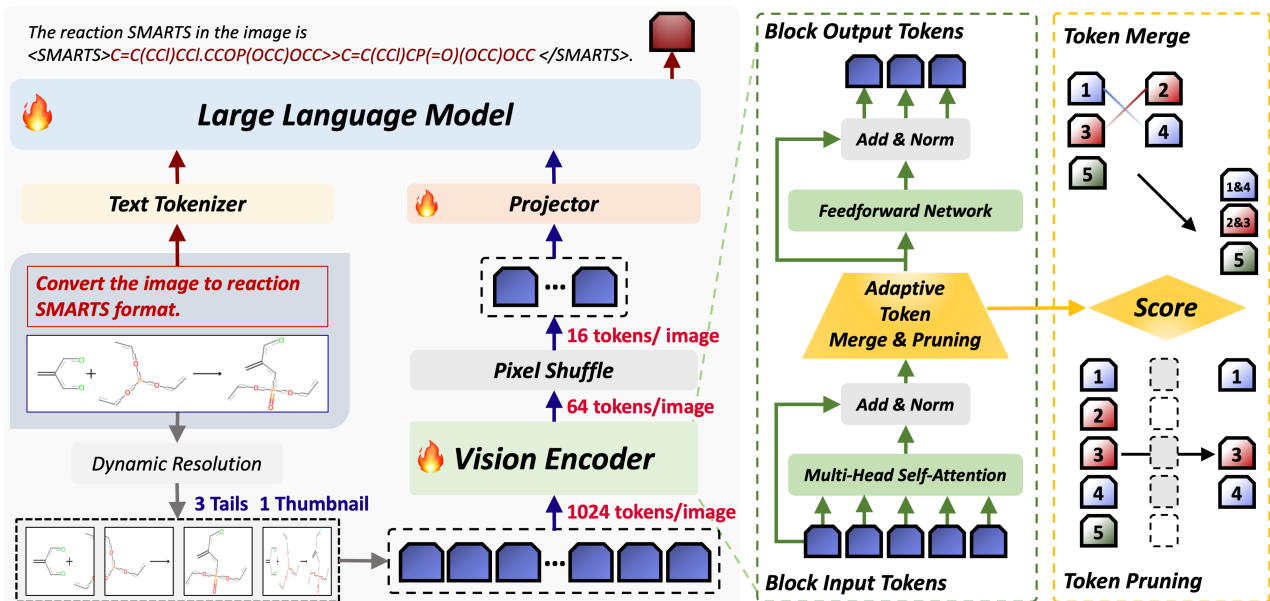


Figure 1: The overview structure of TinyChemVL. For the input image, a dynamic resolution strategy is employed to segment and downsample it to standard resolutions. The vision encoder progressively reduces the number of visual tokens through the adaptive token merging and pruning method. The pixel shuffle operation is the default strategy utilized in InternVL2.5.

a severe modality imbalance when representing chemical molecular images. A standard-resolution ( $448 \times 448$ ) molecular image is encoded into 256 visual tokens, approximately 8.5 times more than the around 30 textual tokens required for its corresponding SMILES representation. (Based on an average of 200k SMILES embedded with the InternVL2.5 tokenizer). As molecular images are characterized by inherent sparsity, most visual tokens correspond to blank backgrounds. This results in significant visual token redundancy, imposing a substantial computational burden on chemical VLMs. Furthermore, the patching process disrupts molecular structures by fragmenting the visual representation. To mitigate the above challenges, we introduce an adaptive token merging and pruning strategy.

### Adaptive Token Merge and Pruning

The visual token reduction strategy, positioned between the attention and FFN layers in each transformer block, adaptively prunes inattentive tokens and merges duplicative ones based on the current visual token distribution (Bolya et al. 2022; Kim et al. 2024). In detail, token scores are first calculated to quantify the task-specific importance of each token. While numerous non-parametric token scoring mechanisms have been proposed (Liang et al. 2022; Xu et al. 2022; Zeng et al. 2025), we adopt the Adaptive Token Sampler (ATS) (Fayyaz et al. 2022) specifically for its ability to incorporate the value matrix  $V$  in calculations.

$$A = \text{Softmax} \left( \frac{Q \cdot K^T}{\sqrt{d}} \right) \quad (1)$$

$$\text{Score}_i = \frac{A_{1,i+1} \times \|V_{i+1}\|}{\sum_{j=1}^N A_{1,j+1} \times \|V_{j+1}\|}, \quad (2)$$

where  $Q$ ,  $K$ ,  $V$  are the query, key, values of self-attention.  $A$  is the attention metric.  $\text{Score}_i$  is the importance score assigned to image tokens.

**Token Pruning** For token pruning, we adopt the standard Top-K selection strategy (Wu et al. 2023; Liang et al. 2022), which preserves a fixed number (K) of the highest-scoring tokens. Since the total number of input tokens at each layer is pre-defined, this approach is equivalent to pruning a fixed number of tokens, which provides predictable control over the compression ratio.

**Token Merge** For token merging, we leverage the Bipartite Soft Matching (BSM) algorithm (Bolya et al. 2022) due to its superior performance. BSM first partitions the tokens into two equal-sized sets. It then constructs a bipartite graph by creating an edge between each token in one set and its most similar counterpart in the other, using cosine similarity as the metric. After selecting the Top-K edges with the highest similarity scores, the remaining connected tokens are merged via a weighted average of their features. Crucially, this merging process is not constrained by spatial adjacency. Non-adjacent tokens can also be merged if they belong to different subsets and are similar enough.

**Proportional Attention** As the sequence of tokens is condensed by merging and pruning after each attention block, we follow ToMe (Bolya et al. 2022) to maintain a row vector  $s$  that tracks the number of original tokens each current token represents, thereby preserving information fidelity. This row vector is also used as a weight to reflect its importance in the calculation of the attention matrix

$$A = \text{Softmax} \left( \frac{Q \cdot K^T}{\sqrt{d}} + \log s \right). \quad (3)$$

**Adaptive Policy** Consistent with trends observed in prior

work (Wu et al. 2023), our analysis of the InternViT-300M-448px-V2.5 encoder confirms that the variance of token significance scores for a given sample correlates positively with model depth. This inspired us to employ an adaptive strategy to mitigate the token redundancy. When token significance scores exhibit low variance, it indicates a convergence of importance, often corresponding to a large background area. In this situation, these uninformative tokens are pruned. Conversely, when the scores show high variance, it signifies that some tokens are much more important than others and likely represent complex structures. In these cases, tokens corresponding to adjacent or functionally similar molecular structures are merged. This approach reduces the number of visual features by removing irrelevant tokens, while simultaneously creating a more efficient and holistic structural representation of the molecule. Furthermore, since the variance of token significance scores differs across instances even within the same layer, we propose a strategy that adapts at both the instance and layer levels.

$$S_{op_i} = \text{var}(\text{Score}_i), \quad (4)$$

$$\text{Policy}_i = \begin{cases} \text{Token Pruning,} & \text{if } S_{op_i} \leq \tau \\ \text{Token Merge,} & \text{otherwise} \end{cases} \quad (5)$$

where  $\tau$  is a decision threshold hyperparameter that balances the trade-off between the pruning and merge strategies, which we set to a default value of  $1e - 5$  based on a statistical analysis of score variances across existing chemical image datasets.

## Data Construction

### Training Data

A primary challenge in training effective VLMs for chemistry is the curation of suitable training datasets. While numerous chemical datasets have been collected, the majority of them are in text (Zhang et al. 2024a; Fang et al. 2023) or graph (Park et al. 2024; Lee et al. 2025) formats, with a notable lack of image data. However, training high-performing VLMs necessitates a large-scale and diverse image-text corpus. As the training corpus of ChemVLM (Li et al. 2025) is not open-source, we leverage the training corpora provided by the more recent work ChemMLLM (Tan et al. 2025) as the initial datasets. However, we find that the scope and diversity of this training dataset are still limited. Therefore, we construct a large-scale training dataset containing both molecular and reaction-level tasks.

**Molecule Recognition** For vision-based chemical tasks, where molecules are represented as images, the capacity to recognize their structures is a foundational skill. This task requires the VLM to parse the visual representation of a molecule into its SMILES format, thereby unlocking subsequent chemical analyses. We collect training corpus from ChEBI-20-MM (Liu, Tao, and Ren 2025), MolGrapher (Morin et al. 2023), MolScribe (Qian et al. 2023), ORDERly (Wigh et al. 2024) and EDU-CHEMC (Chang et al. 2025). For data already available as `<image, SMILES>` pairs, we incorporate them directly into our training corpus. For all other instances, which consist solely of SMILES strings, we

generate the corresponding molecular images using RDKit (Bento et al. 2020) and Indigo (EPAM Systems 2024). To ensure visual diversity, we manually define a diverse range of seed drawing parameters, including color, size and style, which are then randomly combined and applied during image rendering. Furthermore, our experiments reveal that recognizing SMILES strings from handwritten molecular images poses a significant challenge. To address this, we adapt the EDU-CHEMC dataset (Chang et al. 2025) by developing a parser code that converts its native structure-specific markup language (SSML) (Hu et al. 2023) into corresponding SMILES, yielding a paired dataset of handwritten images and their corresponding SMILES string.

**Reaction Recognition.** Beyond individual molecules, we also apply TinyChemVL to the task of whole reaction recognition, which is foundational to understanding chemical reactions. To address the scarcity of public data on chemical reactions, we utilize the same method denoted in molecular image generation and render reaction images from the ORDERly dataset (Wigh et al. 2024). For each input image, the task is to parse the depicted chemical reaction into the `reactants>reagents>solvents>products` format. Each component, expressed in SMILES format, may represent more than one molecule. In cases with multiple molecules, their respective SMILES strings are concatenated and separated by a period (e.g., `c1ccccc1.Cl`).

**Property Prediction.** Following the property prediction task of ChemMLLM, we construct training data involving first decomposing reactions in the ORDERly dataset into molecules, then sampling SMILES based on their length distribution, and finally utilizing RDKit to calculate molecular properties. The properties include molecular weight (MW), logarithm of the Partition Coefficient of a solute between octanol and water (LogP), Topological Polar Surface Area (TPSA), Hydrogen Bond Donor (HBD), Hydrogen Bond Acceptor (HBA), Rotatable Bond (RB), and Quantitative Estimate of Drug-likeness (QED).

**Reaction Prediction.** Drawing inspiration from chemical experts who can predict reaction products simply by observing the structures of the reactants, we construct a vision-based reaction prediction dataset to train TinyChemVL for this task. Specifically, the model learns to predict the product from an image depicting the reactants, reagents, and solvents. Additionally, to better assist the model in understanding the input image, we also construct the training corpus that pairs the input images with the corresponding SMILES strings. All the data for this task is sampled from the ORDERly dataset as well.

**Molecular Image Generation.** ChemMLLM introduces the task of conditional chemical image generation, which requires models to generate a target molecular image based on specified properties or optimization requirements. However, this task remains text-centric, as performance evaluation ultimately relies on converting the generated image back into a SMILES string using tools like MolScribe (Qian et al. 2023). In contrast, we propose a more direct and verifiable paradigm that utilizes executable Python code to render the target molecular image. To implement this, we sample and reconstruct the ChemMLLM training data by replacing

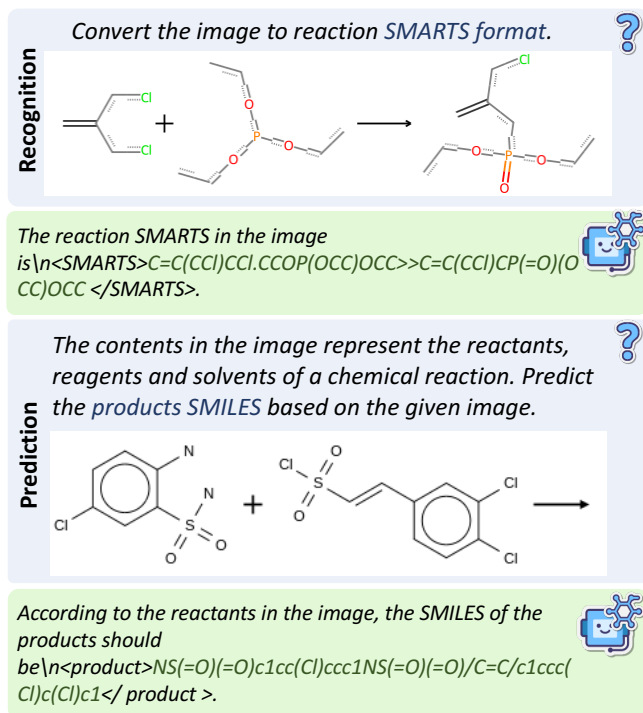


Figure 2: The demonstration about vision-based reaction recognition and prediction tasks in ChemRxn-V.

Task	Total	Ours
Molecule Recognition	500k	385k
Reaction Recognition	200k	200k
Property Prediction	150k	55k
Reaction Prediction	200k	200k
Molecular Image Generation	200k	55k

Table 1: Quantity of training samples across different tasks, showing the total number of samples (**Total**) and the subset generated by ourselves (**Ours**). The remaining samples are collected from existing vision-based chemical datasets.

its `<|image|>` tokens with the corresponding rendering code. We also collect further property-to-image data to expand the training corpus.

Table 1 demonstrates the composition of our training dataset for each task, distinguishing between pre-existing samples and those we construct.

## Evaluation

For the vision-based reaction recognition and prediction tasks, we introduce a new evaluation benchmark named (**ChemRxn-V**). This benchmark, containing reaction recognition and prediction tasks, is curated by filtering the OR-Orderly test set (Wigh et al. 2024) and subsequently rendering the corresponding reaction and reactant images using RDKit and Indigo. Specifically, each task in the ChemRxn-V benchmark contains 5,000 test samples, which are se-

lected based on reaction length to ensure the distribution is representative of the overall dataset. Prediction performance is assessed by computing the RDKit fingerprint similarity between predicted and ground-truth products, following the Mol-instructions (Fang et al. 2023) protocol. In contrast, evaluating recognition performance requires a more detailed approach, as RDKit cannot directly compare reactions. Therefore, we first decompose each reaction into its three components (reactants, reagents/solvents, products), then calculate the fingerprint similarity for each part, and finally compute a weighted average score based on the molecule count per component.

## Training Details

We utilize InternVL2.5-4B as the backbone model for supervised fine-tuning (SFT). The visual token reduction strategy provides the computational efficiency necessary for full-parameter fine-tuning. The model is trained on 8 NVIDIA A100×80G GPUs for 1.5 epochs. The per-device batch size and the gradient accumulation steps are set to 16 and 2, respectively. All the training process is conducted on ms-swift.

## Experiments

### Experiment Settings

**Baselines** To assess the effectiveness of our proposed TinyChemVL, we compare it with existing models in three setups: (1) General-domain open-source VLMs, including InternVL2.5-4B (Chen et al. 2024), Qwen2.5-VL-7B (Bai et al. 2025), LLaVA-v1.5-7B (Liu et al. 2023) and Phi-3.5-vision (Abdin et al. 2024). (2) Proprietary models include GPT-4V (OpenAI 2023) and GPT-4o (OpenAI 2024). (3) Chemical-domain VLMs including ChemVLM (8B, 26B) (Li et al. 2025), ChemDFM-X (13B) (Zhao et al. 2024a) and ChemMLLM (Tan et al. 2025).

**Benchmarks** To ensure a fair and direct comparison with prior works, we align our evaluation protocols with established methods. For the molecule recognition task, we adopt the benchmarks from ChemVLM (ChemOCR) and ChemMLLM (img2smiles), reporting two primary metrics calculated via RDKit, the average Tanimoto similarity (Avg. Sim.) and the Tanimoto hit at 1.0 (Tani@1.0). For property prediction, we also utilize benchmarks in ChemMLLM (img2property) using the Mean Squared Error (MSE) between predicted and ground-truth values for evaluation. For molecular image generation, we assess performance on the two distinct subtasks proposed in ChemMLLM. For the image-to-image (img2img) task, property improvement is quantified by the increase in LogP values between the input and output molecules. For the property-to-image (property2img) task, we calculate the Mean Squared Error (MSE) between the output and the ground-truth input utilizing RDKit. Finally, for reaction-level tasks, we evaluate performance using our proposed ChemRxn-V benchmark, utilizing the aforementioned method for evaluation. To evaluate our adaptive visual token reduction strategy, we assess the computational efficiency of TinyChemVL by comparing its inference throughput and training time against

Models	ChemOCR (%)		img2smiles(%)	
	Avg Sim.	Tani@1.0	Avg Sim.	Tani@1.0
GPT-4V	15.0	2.1	9.74	0.01
GPT-4o	36.8	3.4	29.0	0.01
Qwen2.5-VL-7B	25.5	0.4	28.2	0.03
LLaVA-v1.5-7B	9.0	0.0	5.0	0.00
InternVL2.5-4B	4.4	0.0	2.0	0.01
Phi-3.5-vision	0.4	0.0	1.2	0.01
ChemVLM-26B	71.0	42.9	47.1	12.1
ChemVLM-8B	<u>81.7</u>	<u>57.7</u>	55.0	15.0
ChemMLLM	-	-	75.0	49.0
ChemDFM-X	70.9	36.5	<b>90.9</b>	<b>77.6</b>
TinyChemVL	<b>91.2</b>	<b>77.4</b>	<u>89.5</u>	<u>75.6</u>

Table 2: Evaluation results of the molecular recognition task. The Avg Sim. and Tani@1.0 are the abbreviations of average tanimoto similarity and tanimoto hit 1.0. The best and second-best performances are indicated in **bold** and underline respectively.

Models	img2property ↓						
	MW	LogP	TPSA	HBD	HBA	RB	QED
GPT-4o	7.6e3	5.7	1.2e3	1.8	6.2	11.9	0.03
Qwen2.5VL-7B	5.3e3	22	3.0e3	2.4	12.7	49.4	0.19
LLaVA-v1.5-7B	3.0e4	5.1	5.0e4	53.3	23.7	39.9	1.0e5
Phi-3.5-vision	4.8e4	14.7	2.3e4	3.3	17.7	1.0e6	276.8
InternVL2.5-4B	2.3e4	9.7	5.4e3	1.3e4	6.0e4	2.0e4	2.3
ChemVLM-8B	1.0e4	4.9	5.0e3	4.2	27.2	45.8	0.24
ChemMLLM	<u>789.7</u>	<b>0.7</b>	<u>152.6</u>	<u>0.18</u>	<u>0.8</u>	<u>1.6</u>	<u>0.008</u>
TinyChemVL	<b>488.0</b>	<u>0.9</u>	<b>71.1</b>	<b>0.09</b>	<b>0.3</b>	<b>0.7</b>	<b>0.003</b>

Table 3: Evaluation results on the property prediction task. We manually evaluate models that are not reported in the ChemMLLM and calculate the Mean Squared Error (MSE) for comparison.

two baselines: the general-domain InternVL-2.5-4B and the chemical-domain ChemVLM-8B.

## Main Results

**Model Performance.** The performance of molecular recognition is denoted in the Table 2. As the foundation task of vision-based chemistry, TinyChemVL achieves SOTA performance on the ChemOCR, significantly surpassing existing VLMs. While ChemDFM-X (13B) slightly outperforms TinyChemVL (4B) on the img2smiles task, our model is significantly more parameter-efficient, achieving nearly equivalent performance. Furthermore, we compare TinyChemVL with specialized SMILES OCR models, for example, Decimer (Rajan, Zielesny, and Steinbeck 2021) reports a 92.6% average similarity and a 77.3% Tanimoto@1.0 on the ChemOCR benchmark. TinyChemVL is the first VLM to deliver competitive results with these specialized models. Unlike specialized models that employ specific loss functions for the image-to-SMILES task, our approach demon-

Models	img2img ↑	property2img ↓		
	Increased LogP	MW	LogP	TPSA
GPT-4o	1.95	<u>7633.2</u>	<u>5.7</u>	<u>1209.2</u>
LLaVA-v1.5-7B	-0.86	> 2e6	1104.4	> 1e5
ChemLLM-7B	-	> 7e4	10.0	7666.0
ChemVLM-8B	0.45	> 2e5	9.0	> 2e4
ChemMLLM	<b>4.2</b>	> 3e4	13.2	1191.6
TinyChemVL	<u>2.2</u>	<b>1620.2</b>	<b>0.9</b>	<b>182.5</b>

Table 4: Evaluation results on the molecule image generation task. TinyChemVL achieves superior performance on property2img tasks.

Models	Recognition		Prediction	
	Avg Sim.	EM	Avg Sim.	Tani@1.0
GPT-4o	19.1	0.1	<u>30.4</u>	<u>1.4</u>
Qwen2.5VL-7B	3.8	0.1	11.9	0.0
Phi-3.5-vision	1.5	0.0	0.8	0.0
InternVL2.5-4B	1.6	0.0	2.7	0.0
ChemVLM-8B	0.6	0.0	4.8	0.0
ChemDFM-X	<u>28.32</u>	<u>3.2</u>	12.7	0.7
TinyChemVL	<b>93.4</b>	<b>67.9</b>	<b>78.9</b>	<b>52.4</b>

Table 5: The evaluation results on our proposed ChemRxn-V. EM is the abbreviation of Exact Match.

strates that general-purpose VLMs have the potential to match their performance. This finding indicates that general VLMs can match the performance of specialist models on specialized tasks without sacrificing their broader and versatile capabilities.

The results for the property prediction tasks are summarized in Table 3. TinyChemVL demonstrates superior performance across most tasks, achieving approximately half the Mean Squared Error (MSE) of the second-best model, ChemMLLM. The only suboptimal performance is the LogP img2img task. The performance on the image generation task is denoted in Table 4. We evaluate baseline models (excluding ChemMLLM) on their direct generation of SMILES strings. In contrast, TinyChemVL generates executable Python code that renders the molecule. For a fair assessment, we extract the SMILES from the generated code for comparison. The results show that TinyChemVL achieve superior performance on the property2img task. In the img2img task, TinyChemVL achieves the second-best performance, surpassed only by the ChemMLLM.

Furthermore, we evaluate model performance on our constructed ChemRxn-V, which contains vision-based reaction recognition and prediction tasks. As shown in Table 5, TinyChemVL achieves best performance, significantly outperforming existing VLMs that struggle with reaction-level tasks. Notably, for the reaction prediction task, no existing models attempt to generate products based solely on visual representations of the reactants. For chemical experts, the

Models	Inference		Training
	Sample/s ( $\uparrow$ )	Avg. Tokens ( $\downarrow$ )	Hours ( $\downarrow$ )
ChemVLM-8B	7.41	896	-
InternVL2.5-4B	9.11	894	47*
TinyChemVL	11.84	108	15

Table 6: The results of inference efficiency on ChemOCR. \* denotes the estimated time by tqdm. The Avg. Tokens denotes the average number of tokens per input, including both visual and textual tokens.

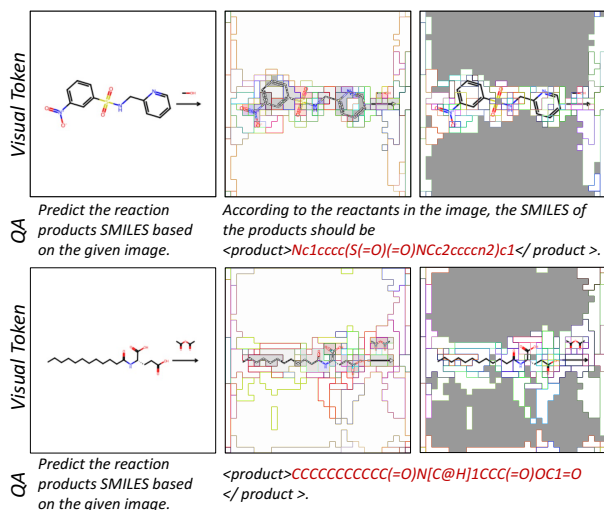


Figure 3: The odd rows (QA) show how the method prunes blank background tokens (indicated by grey areas) and merges molecular structure tokens in a reaction image. The even rows (Visual Tokens) display the associated task instruction and the final answer generated by TinyChemVL.

molecular image is the most intuitive medium for understanding reactions and predicting products. Unlike textual notations such as SMILES, which are essentially a human-defined grammar for computation, molecular images offer a more natural representation. Therefore, enabling models to predict reaction products directly from molecular images represents a crucial research frontier. Success in this area is pivotal for establishing VLMs as indispensable tools for chemical research.

Overall, TinyChemVL achieves SOTA performance across a diverse suite of vision-based chemical tasks at both the molecular and reaction levels. Furthermore, we propose ChemRxn-V, a vision-based benchmark for reaction-level tasks that poses a significant challenge to existing models, thereby expanding the scope of this research domain.

**Model Efficiency.** We conduct two primary comparisons to evaluate our proposed TinyChemVL. First, to specifically assess the effectiveness of our adaptive visual token reduction method, we benchmark TinyChemVL against its backbone model, InternVL2.5-4B. Second, to evaluate the overall efficiency of our approach relative to specialized models,

Ablation (tokens/image)	ChemOCR	ChemRxn-V	
	Recognition	Recognition	Prediction
16	77.4	62.7	52.4
4	76.2(-1.2)	59.5(-3.2)	50.1(-2.3)

Table 7: The ablation study of different numbers of visual tokens on molecular and reaction tasks. Tani@1.0 and exact match are utilized for comparison.

we compare TinyChemVL with ChemVLM-8B, the well-established VLM in the chemical domain. As all these models are based on the InternVL family, we utilize LMDeploy (Contributors 2023) as the inference toolkit for fairly comparison. We evaluate model efficiency on the ChemOCR dataset by measuring both inference throughput (Samples/s) and computational overhead, which is quantified by the average number of input tokens (Avg. Tokens). Furthermore, we compare the training time cost of TinyChemVL against InternVL-2.5-4B on our training corpus. To ensure a fair comparison under identical hardware constraints, we run each model at the maximum batch size it can accommodate within the GPU memory budget of A100. This setting reveals that TinyChemVL can utilize a batch size four times larger than InternVL-2.5-4B, forming the basis for our subsequent time measurements.

The results are shown in Table 6. TinyChemVL achieves the highest inference speed, a result attributed to its fewer input tokens, which requires only about 1/8 input tokens compared to competing models. Notably, TinyChemVL exhibits a significant reduction in training time compared to InternVL2.5-4B, which enhances the efficiency of finetuning and promotes the iteration process.

## Ablation Study

To demonstrate the efficacy of our visual token reduction strategy, Figure 3 visualizes the compression results on two chemical reaction images. These results show that our method effectively prunes tokens representing solely the blank background while preserving those crucial for representing the molecules, without influencing the answers. Also, we conduct an ablation study on the impact of reduced visual token quantity with the same training data. As shown in Table 7, further reducing visual tokens from 16 to 4 per image results in performance drops, demonstrating the insufficiency of 4 tokens for representing reaction images.

## Conclusion

In this work, we propose TinyChemVL, an efficient and powerful chemical VLM that performs superiorly on vision-based molecular and reaction tasks. By adopting our token reduction strategy, we reduce the number of visual tokens to 1/16th of the original count, ensuring high efficiency for both training and inference. Also, we propose a vision-based reaction-level benchmark, ChemRxn-V, to assess the vision-based reaction recognition and prediction capacities. ChemRxn-V and TinyChemVL could generalize chemical VLMs into more complex scenarios.

## Acknowledgements

This work was supported by the robotic AI-Scientist platform of Chinese Academy of Sciences.

## References

- Abdin, M.; Aneja, J.; Awadalla, H.; Awadallah, A.; Awan, A. A.; Bach, N.; Bahree, A.; Bakhtiari, A.; Bao, J.; Behl, H.; et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Beltagy, I.; Lo, K.; and Cohan, A. 2019. SciBERT: A pre-trained language model for scientific text. *arXiv preprint arXiv:1903.10676*.
- Bento, A. P.; Hersey, A.; Félix, E.; Landrum, G.; Gaulton, A.; Atkinson, F.; Bellis, L. J.; De Veij, M.; and Leach, A. R. 2020. An open source chemical structure curation pipeline using RDKit. *Journal of Cheminformatics*, 12: 1–16.
- Bi, J.; Wang, Y.; Chen, H.; Xiao, X.; Hecker, A.; Tresp, V.; and Ma, Y. 2024. Visual Instruction Tuning with 500x Fewer Parameters through Modality Linear Representation-Steering. *arXiv preprint arXiv:2412.12359*.
- Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2022. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*.
- Chang, Q.; Chen, M.; Pi, C.; Hu, P.; Zhang, Z.; Ma, J.; Du, J.; Yin, B.; and Hu, J. 2025. RFL: Simplifying Chemical Structure Recognition with Ring-Free Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 2007–2015.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Contributors, L. 2023. LMDeploy: A Toolkit for Compressing, Deploying, and Serving LLM. <https://github.com/InternLM/lmdeploy>.
- Edwards, C.; Lai, T.; Ros, K.; Honke, G.; Cho, K.; and Ji, H. 2022. Translation between molecules and natural language. *arXiv preprint arXiv:2204.11817*.
- Edwards, C.; Zhai, C.; and Ji, H. 2021. Text2mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 595–607.
- EPAM Systems. 2024. Indigo Toolkit. <https://github.com/epam/Indigo>. Accessed: June 20, 2025.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Fang, Y.; Liang, X.; Zhang, N.; Liu, K.; Huang, R.; Chen, Z.; Fan, X.; and Chen, H. 2023. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. *arXiv preprint arXiv:2306.08018*.
- Fayyaz, M.; Koochpayegani, S. A.; Jafari, F. R.; Sengupta, S.; Joze, H. R. V.; Sommerlade, E.; Pirsiavash, H.; and Gall, J. 2022. Adaptive token sampling for efficient vision transformers. In *European Conference on Computer Vision*, 396–414. Springer.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hu, J.; Wu, H.; Chen, M.; Liu, C.; Wu, J.; Yin, S.; Yin, B.; Yin, B.; Liu, C.; Du, J.; et al. 2023. Handwritten chemical structure image to structure-specific markup using random conditional guided decoder. In *Proceedings of the 31st ACM International Conference on Multimedia*, 8114–8124.
- Kim, M.; Gao, S.; Hsu, Y.-C.; Shen, Y.; and Jin, H. 2024. Token fusion: Bridging the gap between token pruning and token merging. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1383–1392.
- Krenn, M.; Ai, Q.; Barthel, S.; Carson, N.; Frei, A.; Frey, N. C.; Friederich, P.; Gaudin, T.; Gayle, A. A.; Jablonka, K. M.; et al. 2022. SELFIES and the future of molecular string representations. *Patterns*, 3(10).
- Lee, C.; Song, Y.; Jeong, Y.; Ko, H.; Hormazabal, R.; Han, S.; Bae, K.; Lim, S.; and Kim, S. 2025. Mol-LLM: Generalist Molecular LLM with Improved Graph Utilization. *arXiv preprint arXiv:2502.02810*.
- Li, J.; Zhang, D.; Wang, X.; Hao, Z.; Lei, J.; Tan, Q.; Zhou, C.; Liu, W.; Yang, Y.; Xiong, X.; et al. 2025. Chemvlm: Exploring the power of multimodal large language models in chemistry area. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 415–423.
- Liang, Y.; Ge, C.; Tong, Z.; Song, Y.; Wang, J.; and Xie, P. 2022. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800*.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning.
- Liu, P.; Tao, J.; and Ren, Z. 2025. A quantitative analysis of knowledge-learning preferences in large language models in molecular science. *Nature Machine Intelligence*, 1–13.
- Lu, P.; Qiu, L.; Chang, K.-W.; Wu, Y. N.; Zhu, S.-C.; Rajpurohit, T.; Clark, P.; and Kalyan, A. 2022. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv preprint arXiv:2209.14610*.
- M. Bran, A.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; and Schwaller, P. 2024. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5): 525–535.

- Morin, L.; Danelljan, M.; Agea, M. I.; Nassar, A.; Weber, V.; Meijer, I.; Staar, P.; and Yu, F. 2023. MolGrapher: graph-based visual recognition of chemical structures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19552–19561.
- OpenAI. 2023. GPT-4V(ision) System Card.
- OpenAI. 2024. GPT-4o. Accessed: 2024-05-13.
- Park, J.; Bae, M.; Ko, D.; and Kim, H. J. 2024. LLaMo: Large Language Model-based Molecular Graph Assistant. *arXiv preprint arXiv:2411.00871*.
- Qian, Y.; Guo, J.; Tu, Z.; Li, Z.; Coley, C. W.; and Barzilay, R. 2023. MolScribe: robust molecular structure recognition with image-to-graph generation. *Journal of Chemical Information and Modeling*, 63(7): 1925–1934.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Rajan, K.; Zielesny, A.; and Steinbeck, C. 2021. DECIMER 1.0: deep learning for chemical image recognition using transformers. *Journal of Cheminformatics*, 13: 1–16.
- Tan, Q.; Zhou, D.; Xia, P.; Liu, W.; Ouyang, W.; Bai, L.; Li, Y.; and Fu, T. 2025. ChemMLLM: Chemical Multimodal Large Language Model. *arXiv preprint arXiv:2505.16326*.
- Wigh, D. S.; Arrowsmith, J.; Pomberger, A.; Felton, K. C.; and Lapkin, A. A. 2024. Orderly: data sets and benchmarks for chemical reaction data. *Journal of Chemical Information and Modeling*, 64(9): 3790–3798.
- Wu, X.; Zeng, F.; Wang, X.; and Chen, X. 2023. Ppt: Token pruning and pooling for efficient vision transformers. *arXiv preprint arXiv:2310.01812*.
- Xu, Y.; Zhang, Z.; Zhang, M.; Sheng, K.; Li, K.; Dong, W.; Zhang, L.; Xu, C.; and Sun, X. 2022. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 2964–2972.
- Zeng, F.; and Yu, D. 2024. M2M-TAG: Training-Free Many-to-Many Token Aggregation for Vision Transformer Acceleration. In *Workshop on Machine Learning and Compression, NeurIPS 2024*.
- Zeng, F.; Yu, D.; Kong, Z.; and Tang, H. 2025. Token Transforming: A Unified and Training-Free Token Compression Framework for Vision Transformer Acceleration. *arXiv preprint arXiv:2506.05709*.
- Zhang, D.; Liu, W.; Tan, Q.; Chen, J.; Yan, H.; Yan, Y.; Li, J.; Huang, W.; Yue, X.; Ouyang, W.; et al. 2024a. Chemllm: A chemical large language model. *arXiv preprint arXiv:2402.06852*.
- Zhang, D.; Yu, Y.; Dong, J.; Li, C.; Su, D.; Chu, C.; and Yu, D. 2024b. Mm-llms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*.
- Zhang, J.; Bian, Y.; Chen, Y.; and Yao, Q. 2024c. Unimot: Unified molecule-text language model with discrete token representation. *arXiv preprint arXiv:2408.00863*.
- Zhang, P.; Dong, X.; Zang, Y.; Cao, Y.; Qian, R.; Chen, L.; Guo, Q.; Duan, H.; Wang, B.; Ouyang, L.; et al. 2024d. Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output. *arXiv preprint arXiv:2407.03320*.
- Zhang, Q.; Ding, K.; Lv, T.; Wang, X.; Yin, Q.; Zhang, Y.; Yu, J.; Wang, Y.; Li, X.; Xiang, Z.; et al. 2025. Scientific large language models: A survey on biological & chemical domains. *ACM Computing Surveys*, 57(6): 1–38.
- Zhao, X.; Jiang, D.; Zeng, Z.; Chen, L.; Qiu, H.; Huang, J.; Zhong, Y.; Zheng, L.; Cao, Y.; and Ma, L. 2025a. Vinci-Coder: Unifying Multimodal Code Generation via Coarse-to-fine Visual Reinforcement Learning. *arXiv preprint arXiv:2511.00391*.
- Zhao, X.; Luo, X.; Shi, Q.; Chen, C.; Wang, S.; Liu, Z.; and Sun, M. 2025b. Chartcoder: Advancing multimodal large language model for chart-to-code generation. *arXiv preprint arXiv:2501.06598*.
- Zhao, Z.; Chen, B.; Li, J.; Chen, L.; Wen, L.; Wang, P.; Zhu, Z.; Zhang, D.; Li, Y.; Dai, Z.; et al. 2024a. ChemDFM-X: towards large multimodal model for chemistry. *Science China Information Sciences*, 67(12): 1–2.
- Zhao, Z.; Ma, D.; Chen, L.; Sun, L.; Li, Z.; Xu, H.; Zhu, Z.; Zhu, S.; Fan, S.; Shen, G.; et al. 2024b. Chemdfm: Dialogue foundation model for chemistry. *arXiv e-prints*, arXiv–2401.