CheMM-R1: Enhancing Chemical Structure Recognition and Elucidation with Reasoning Multimodal Large Language Models

Liting Huang[1]*, **Zhihao Zhang**[2]*, **Shoujin Wang**[1]
University of Technology Sydney[1], University of New South Wales[2]
liting.huang@student.uts.edu.au
zhihao.zhang11@unsw.edu.au
shoujin.wang@uts.edu.au

Abstract

While Multimodal Large Language Models (MLLMs) demonstrate strong reasoning capabilities, they lack domain-specific expertise to effectively perform chemical tasks. For example, existing MLLMs struggle with both the lower-level task of molecular structure recognition and the higher-level task of chemical spectral data elucidation. When faced with complex molecular structures and multimodal chemical data (including spectral images and texts), they often fail to provide reliable inference, resulting in poor performance. Moreover, there are no benchmark datasets for evaluating multi-step multimodal reasoning capacities in the chemistry domain. To this end, we establish CheMM-Bench, a comprehensive benchmark dataset with 48,500 reasoning steps across four chemical tasks (SmilesQA, IupacQA, MwQA, SpectraQA) for evaluating visual reasoning in both molecular structure recognition and spectral analysis. On top of this, we present CheMM-R1, a state-of-the-art chemistry-specific MLLM trained with CheMMGRPO, a novel adaptation of Group Relative Policy Optimisation tailored for chemical reasoning. CheMMGRPO employs domain-specific reward functions to assess chemical validity, structural accuracy, format compliance, and factual correctness. CheMM-R1 surpasses leading proprietary models (GPT-o3, Gemini-2.5-Pro, Claude-3.5-Sonnet, and Grok-2) across all CheMM-Bench tasks. The evaluation code and model are publicly available [1].

1 Introduction

Recent progressive multimodal large language models (MLLMs) shows substantial advances in reasoning tasks (Hurst et al., 2024), which process and integrate information from multiple modalities such as text, images, and structured data. These methods often rely on supervised fine-tuning (SFT) with large-scale, high-quality Chain-of-Thought (CoT) data generated by advanced models such as Gemini-2.5-Pro, enabling MLLMs to acquire autonomous reasoning capabilities (Yao et al., 2024). Furthermore, advanced reinforcement learning techniques, such as GRPO (Guo et al., 2025) encourages LLMs to generate high-reward reasoning paths through reward functions to enable model's reasoning capabilities. With these approaches applied to MLLMs (Zhang et al., 2025) and their ability to handle diverse data types, MLLMs possess high potential in scientific contexts and have found broad application in specialised domains like chemistry (Li et al., 2025). Since chemical information is represented across diverse formats that provide complementary structural insights (Klein, 2013), it inherently involves multimodal data such as visual molecular structures, textual nomenclature (SMILES (Weininger, 1988), IUPAC names (Favre and Powell, 2013)), and spectroscopic pattern plots (Nuclear Magnetic Resonance (NMR), Infrared (IR), mass spectra). Chemists typically spend significant time converting manually drawn chemical structures into machine-readable formats and interpreting spectral data for molecular structure determination. Although chemists often possess intuitive understanding of synthesised molecules, verification requires integrating multiple spectroscopic techniques, as each provides distinct structural information. However, manual interpretation of spectroscopic data is labor-intensive and highly subjective, leading to potential inconsistencies.

Despite this clear need for automated multimodal chemical reasoning, current approaches face significant limitations. Traditional Chemical OCR models (Qian et al., 2023) and chemistry-specific MLLMs (Li et al., 2025) have achieved notable success in converting chemical images to SMILES, yet their reasoning abilities remain constrained when handling complex molecular structures. More critically, many existing models are limited by re-

*Equal contribution
[1] https://github.com/liting980713/CheMM-R1

Findings of the Association for Computational Linguistics: ACL 2026, pages 26902–26921

liance on textual modalities (Alberts et al., 2024b), making them unable to process the visual spectral graphs chemists use. Others support only single spectral types (Alberts et al., 2024a; Huang et al., 2021), restricting comprehensive molecular structure inference. While recent reasoning MLLMs excel at general or mathematical tasks (Zhang et al., 2025), they lack domain-specific chemical expertise, resulting in limited generalisation to chemistry-specific problems requiring specialised knowledge of bonding and spectroscopic interpretation. Bridging this gap requires a chemistry-specific MLLM with advanced reasoning capabilities. By integrating textual and visual modalities across molecular structures and spectroscopic data, such a model can provide reliable reasoning and problem-solving, enhancing analytical efficiency while overcoming current limitations in generalisation within the chemical domain. However, achieving this goal faces two key challenges: first, the absence of benchmarks for multimodal reasoning tasks in chemistry; and second, the lack of multimodal language models capable of integrating chemical domain knowledge to support reasoning.

To address theses problems, we first introduce CheMM-Bench, a benchmark specifically designed to assess multi-step visual reasoning tasks in chemistry. The benchmark comprises 48,500 reasoning steps across four datasets: SmilesQA, IupacQA, MwQA, and SpectraQA. Each dataset targets a specific aspect of model capability, enabling comprehensive evaluation. SmilesQA, IupacQA, and MwQA focus on tasks related to molecular structure recognition. These tasks evaluate the ability to derive SMILES notation from molecular images, recognize IUPAC names, and identify molecular masses, collectively assessing both OCR performance and reasoning within the chemical domain. SpectraQA extends evaluation to spectroscopic interpretation, requiring integration of multiple spectral data types for structure determination, mirroring real-world workflows. Furthermore, to evaluate the effectiveness of CheMM-Benc and the MMLMs capable of integrating chemical domain knowledge to support reasoning., we introduce CheMM-R1, a state-of-the-art (SOTA) Multimodal Large Language Model specifically designed for the chemical domain. Inspired by previous studies' (Wei et al., 2025) approach of cold-start and reinforcement learning, we develop multimodal Chain-of-Thought (CoT) patterns for chemistry knowledge injection and utilise Group Relative Policy Optimization (GRPO) (Guo et al., 2025) for reinforcement learning. Through this approach, CheMM-R1 surpasses the performance of current proprietary models, including GPT-O3, Gemini-2.5-Pro, Claude-4-Sonnet, and Grok-4 for all tasks in CheMM-Bench. Our key contributions are:

- **CheMM-Bench.** We create a stepwise visual reasoning benchmark for chemistry, containing 48,500 ultra-long reasoning steps with multiple images across four datasets assessing distinct model capabilities.
- **CheMMGRPO and CheMM-R1.** We propose CheMMGRPO, a chemical knowledge-integrated method that improves MLLM visual reasoning in chemistry. With the proposed CheMMGRPO, we develop CheMM-R1, a chemistry-focused MLLM with superior reasoning capabilities and support for long texts and multiple images.
- **Experimental comparison.** We evaluate CheMM-R1 against existing MLLMs, showing that current reasoning models including GPT-O3 and Gemini-2.5-Pro struggle to accurately understand chemical structures and provide reliable compound information. CheMM-R1 achieves SOTA performance with significant margins over all existing models.

2 Related Work

2.1 Reasoning MLLMs

Inspired by advances in CoT reasoning within LLMs, MLLMs have experienced substantial improvements in reasoning tasks (Hurst et al., 2024; Wu et al., 2024). Current MLLM reasoning methods typically rely on advanced models such as GPT-4 to generate extensive high-quality CoT data, subsequently employing SFT to enhance autonomous reasoning capabilities (Yao et al., 2024). Traditional approaches, including Best-of-N, Beam Search, and Monte Carlo Tree Search (Xie et al., 2024), employ complex search mechanisms to generate extensive reasoning data and apply SFT to enable MLLMs to explore CoT for complex problem solving (Huang et al., 2025). Driven by developments in reinforcement learning, recent works attempt to enhance MLLM reasoning capabilities through reinforcement learning (Zhang et al., 2025). However, these works typically focus on general

Benchmark	Text	Image	Multi-Image	Reasoning Steps	Spectral Data
ChemOCR	✓	✓	×	×	×
MMCR-Bench	✓	✓	×	×	×
MMChemBench	✓	✓	×	×	×
ChemVLM	✓	✓	×	×	×
ChemDFM-X	✓	✓	×	×	✓(MS, IR only)
MolPuzzle	✓	×	×	×	×
CheMM-Bench (Ours)	✓	✓	✓	✓(48,500 steps)	✓(5 types)

Table 1: Comparison of CheMM-Bench with existing multimodal chemistry benchmarks.

and mathematical data types, exhibiting limited generalisation to specific domains like chemistry.

2.2 Chemistry MLLMs

MLLMs have demonstrated significant potential in the chemical domain (Li et al., 2025). Previous work has focused on extracting visual information from chemical images for downstream chemical tasks. Molecular structure recognition has been explored in models such as ChemVLM (Li et al., 2025) and ChemDFM-X (Zhao et al., 2024). While these benchmarks have made important contributions, they primarily evaluate direct answer generation. Without explicit reasoning steps, models learn pattern matching rather than systematic chemical analysis. Furthermore These models are primarily designed to convert chemical images into SMILES representations. Only a few try to utilise spectral data for molecular structure prediction. ChemDFM-X (Zhao et al., 2024) attempt to predict molecular structures using only mass spectrometry and IR data. While MolPuzzle (Guo et al., 2024) attempts molecular structure elucidation but rely on text-based spectral descriptions rather than visual spectral images, and evaluate sub-tasks independently rather than as an end-to-end integration. Similarly, single-spectrum approaches (e.g., IR-only models) are chemically insufficient for full elucidation, as IR cannot determine complete molecular connectivity or stereochemistry without accompanying NMR and MS data. These models typically depend on similarity searches that can yield plausible outputs, but lack the ability to provide explanatory justification. Consequently, current chemical MLLMs struggle to generate outputs that align with chemical intuition and principles. To address this limitation, our model employs logical aggregation of structural features, with a focus on stepwise inference of molecular connectivity. Specifically, the model employs stepwise reasoning: identifying functional groups, combining them into substructures, and ultimately reconstructing the entire molecule. Such an interpretable, stepwise reasoning process strengthens the model's reliability and trustworthiness. To highlight our contributions, Table 1 compares CheMM-Bench against existing datasets. CheMM-Bench is the first to require simultaneous visual integration of five spectral types while demanding explicit, long-context Chain-of-Thought (CoT) reasoning.

3 CheMM-Bench Dataset

CheMM-Bench is a benchmark dataset designed to evaluate organic chemistry reasoning in multimodal large language models. As shown in Figure 2, we utilise the spectroscopic dataset from Alberts (Alberts et al., 2024b), extracted from the USPTO reaction database (Lowe, 2012), containing molecular structures as SMILES and simulated IR, ^{1}H-NMR, ^{13}C-NMR, positive-ion MS, and negative-ion MS spectra for 794,403 realistic molecules from US patents. We randomly extracted 48,500 unique molecules: 26,500 for molecular structure recognition and 22,000 for structure elucidation. CheMM-Bench comprises four tasks: SmilesQA, IupacQA, and MwQA for structure recognition, and SpectraQA for structure elucidation.

3.1 Structure Recognition Task

SmilesQA. SmilesQA is designed for question answering, where the task is to derive the SMILES representation of each molecule, which encodes its unique chemical structure, from its 2D structure. **IupacQA.** IupacQA is formulated as a question–answering task in which the question refers to a molecule's 2D structure and the answer is its IUPAC name. The IUPAC naming system provides a universally recognised standard, whereby each name uniquely corresponds to a specific chemical compound. We leverage PubChem (Kim et al., 2025) to match each molecule's unique IUPAC name from chemical structures as answers. The derivation process involves generating the SMILES representation from the 2D

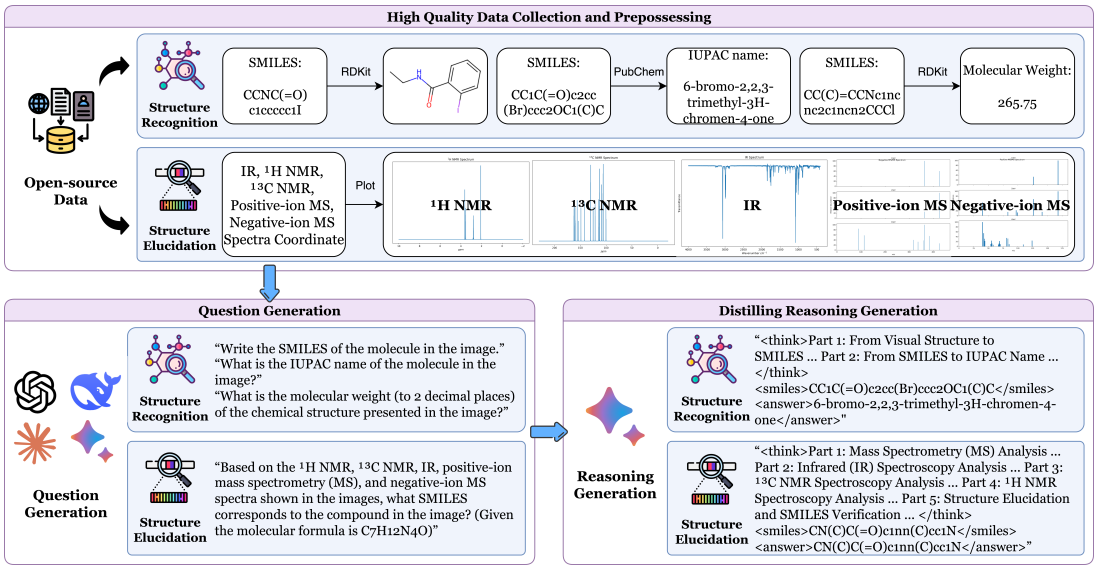

Figure 1: Overview of CheMM-Bench data collection. Structure recognition includes single molecular image; structure elucidation contains six spectral images. Chain-of-Thought reasoning is distilled from Gemini-2.5-Pro.

structure and subsequently translating it into the IUPAC name. **MwQA.** In MwQA, a molecule's 2D structure serves as the input and its molecular weight as the answer. Molecular weights are computed from chemical structures using RDKit (online, 2024b). The inference process involves deriving the SMILES representation from the 2D structure, followed by calculating the molecular weight.

Data Collection. We merge three datasets on multimodal molecular structures to enable extraction of chemical structure (SmilesQA), IUPAC name (IupacQA), and molecular weight (MwQA). The generation process: (1) extract 26,500 unique molecules from Alberts' spectroscopic dataset (Alberts et al., 2024b); (2) convert SMILES to 2D structures using RdKit (online, 2024b); (3) generate QA pairs using four LLMs (DeepSeek-R1 (Guo et al., 2025), ChatGPT-4o-mini (Hurst et al., 2024), Gemini-2.5-Flash (Comanici et al., 2025), and Claude-Sonnet-4 (Anthropic, 2025)), each producing five questions per dataset (20 total), randomly matched with answers. For reasoning generation, we use a semi-automated pipeline with Gemini-2.5-Pro to produce stepwise reasoning in `<think>` tags, SMILES in `<smiles>` tags, and final answers in `<answer>` tags. **Refer to Figures 5-7 in appendix for specific prompts for guided reasoning.**

3.2 Structure Elucidation Task

Nuclear Magnetic Resonance. Nuclear Magnetic Resonance (NMR) spectroscopy detects transitions between nuclear spin states in a magnetic field. Since each isotope absorbs at a specific frequency, NMR can distinguish different elements and typically detects one isotope at a time (e.g., ^{1}H-NMR, ^{13}C-NMR, ^{31}P-NMR). We utilise ^{1}H-NMR and ^{13}C-NMR to probe molecular structures. ^{1}H-NMR detects hydrogen nuclei, revealing the number, type, and chemical environment of hydrogen atoms to identify functional groups and elucidate structures. ^{13}C-NMR identifies carbon atoms through their distinct chemical environments, with each unique carbon producing a characteristic peak whose chemical shift indicates its electronic environment and nearby functional groups. **Infrared.** Similar to NMR, infrared (IR) spectroscopy is among the most widely applied spectroscopic techniques. By recording the absorption of infrared radiation at characteristic wavelengths, it provides insights into functional groups and overall molecular structure. In IR spectroscopy, molecular bonds absorb infrared light, leading to changes in bond length or angle. These absorptions represent as peaks at characteristic wavenumbers, each corresponding to a particular vibrational mode. As different functional groups absorb different frequencies of infrared radiation (e.g., alcohols at 3200–3300 cm^{-1}), the resulting spectra enable chemists to determine the presence of specific bonds or groups and to identify molecular structures. **Mass Spectrometry.** In addition to NMR and IR, mass spectrometry (MS) represents a key analytical technique, wherein the mass-to-charge ratios of either positively charged ions or negatively charged ions are measured to provide insights into molecular mass, elemental composition, and structural features. In mass spec-

trometry, molecules are ionized, the ions are separated using electric and magnetic fields, and detected to produce a mass spectrum, a plot of ion intensity versus mass-to-charge ratio. Each peak in the spectrum reflects the abundance of ions at a particular mass-to-charge ratio. Interpretation of these peaks provides insights into molecular mass, fragmentation patterns, and the identification of the compound. (Klein, 2013) **SpectraQA.** The dataset is formulated as a question answering benchmark, in which the question is defined by five types of molecule's spectral data including IR, ^{1}H-NMR, ^{13}C-NMR, positive-ion MS, and negative-ion MS spectra, and the answer is the corresponding SMILES representation. The task requires inferring the SMILES representation of a molecule directly from its spectral features.

Data Collection. We leverage five different types of spectral data, including IR, ^{1}H-NMR, ^{13}C-NMR, positive-ion MS, and negative-ion MS spectra, to improve the model's ability to predict molecular structures. The data production process is as follows: (1) We randomly extracted 22,000 unique molecules from Alberts' spectroscopic dataset, including SMILES, IR, ^{1}H-NMR, ^{13}C-NMR, positive-ion MS, and negative-ion MS spectra; (2) Different types of spectral data are converted into graphs as image inputs; (3) We utilise the same four LLMs employed in the structure recognition stage to generate 20 questions. The questions are randomly matched with corresponding answers to construct QA pairs. The reasoning generation follows the same method as structure recognition. **We develop specific prompts for structure elucidation as shown in Appendix Figure 8.**

3.3 Data Statistics and Analysis

Figure 2 (A) illustrates the distribution of SMILES lengths, with structure recognition exhibiting a broader distribution, while structure elucidation is shifted toward shorter lengths. The distribution for structure recognition is centred around a SMILES length of 40–50, whereas the peak distribution for structure elucidation occurs at approximately 25–35. Consequently, the structure recognition task typically involves longer and more complex molecules, whereas the structure elucidation task focuses on shorter and simpler SMILES strings. The distribution of heavy atom counts in molecules for the structure recognition and structure elucidation tasks is presented in Figure 2 (B). The peak for structure elucidation occurs around 16–20 heavy atoms, indicating that this task is biased toward smaller to medium-sized molecules. Conversely, the peak for structure recognition occurs at approximately 25–30 heavy atoms, reflecting a tendency toward more complex molecular structures. Furthermore, the chemical similarity of all molecules in the structure recognition and structure elucidation tasks is assessed using the Tanimoto similarity of their chemical fingerprints, as shown in Figure 2 (C). Both tasks display a similar distribution, with a clear peak around 0.1. This indicates that most molecular pairs in both tasks have relatively low similarity, reflecting high structural diversity. Additionally, Figure 2 (D) depicts the distribution of top 20 functional groups present in molecules for the structure recognition and structure elucidation tasks. This figure indicates that Alkanes, Arenes, and Ethers represent the most prevalent functional groups across both tasks. **See Appendix A for an overview of the number of instances in the dataset. Refer to Appendix B for a detailed functional group analysis.**

4 CheMM-R1 Method and Experiment

Equipping Multimodal Large Language Models (MLLMs) with domain-specific knowledge and reasoning abilities typically requires both knowledge injection (Mecklenburg et al., 2024) and reinforcement learning (Wei et al., 2025). Therefore, we propose CheMMGRPO, a novel reinforcement fine-tuning framework designed for multimodal chemistry understanding that enables self-improvement in reasoning. This framework consists of two stages: a cold-start phase, where we incorporate organic chemistry knowledge and construct Chain-of-Thought (Wei et al., 2022) reasoning patterns from collected data, followed by a reinforcement learning phase using online policy optimisation to further enhance the model's reasoning capabilities.

4.1 CheMMGRPO with Cold Start

To prevent the early stable of RL training from the base model without specific organic chemistry knowledge, for CheMM-R1 we utilise the multimodal data $D_s = \{q^n, \tau^n\}_{n=1}^{N}$ from CheMM-Bench with CoT reasoning, where each data includes a specific molecule or spectra question q and the reasoning path τ. We fine-tune the base model on the cold-start data and then apply a multimodal reinforcement learning algorithm based on GRPO (Guo et al., 2025), to facilitate model rea-

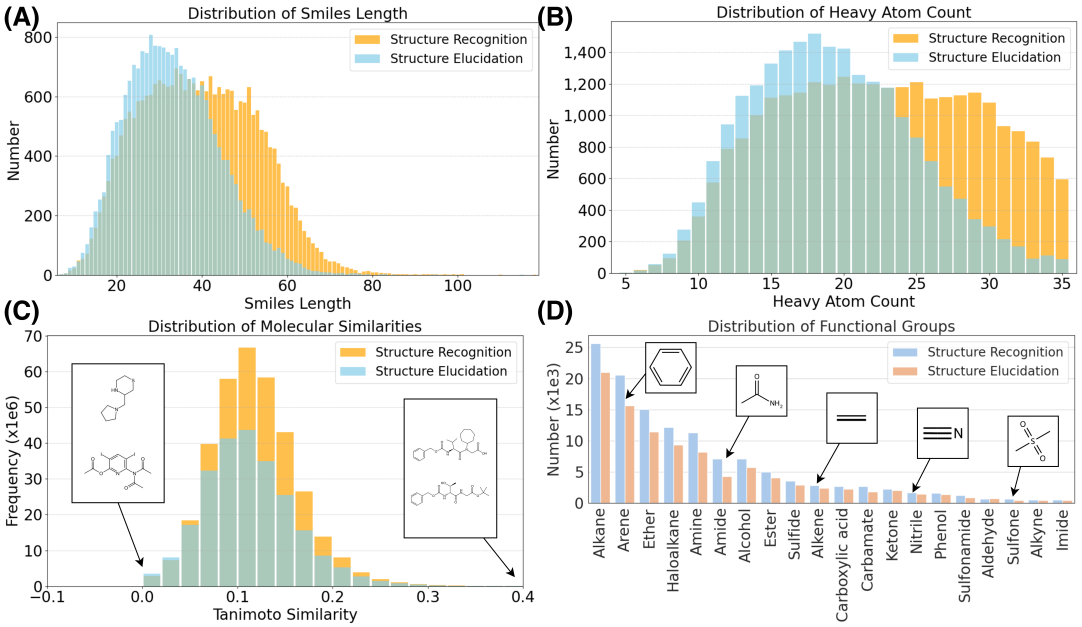

Figure 2: Overall statistics of CheMM-Bench. (A) Distribution of SMILES length; (B) Distribution of heavy atom count; (C) Distribution of molecular similarity; (D) Distribution of top 20 functional groups.

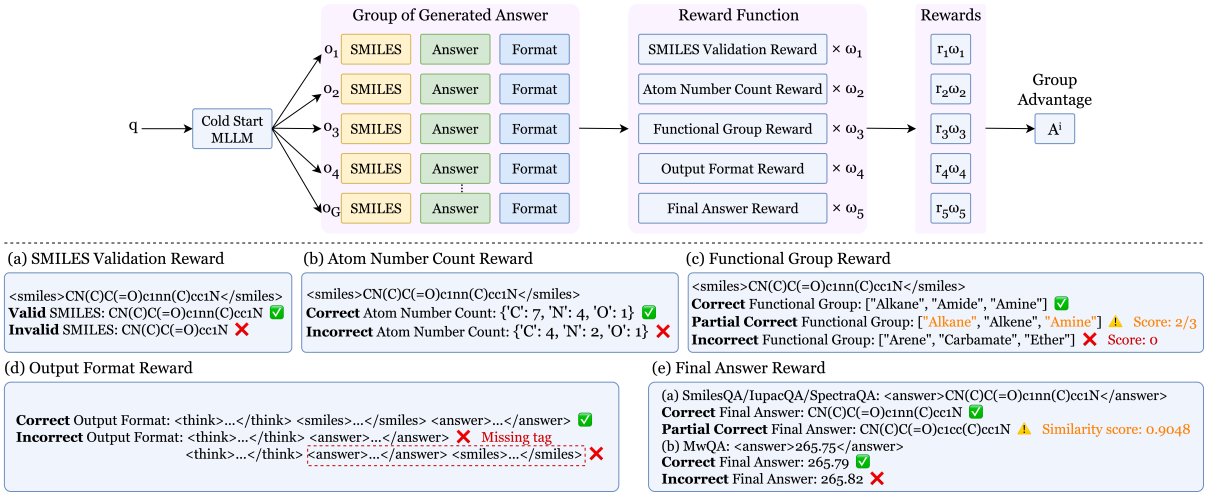

Figure 3: Overview of the proposed CheMMGRPO with chemistry-specific reward functions.

soning and persist chemistry knowledge during the group generation of reasoning trajectories.

For each input question q, the optimisation begins by sampling a group of G distinct outputs $\{o_1, o_2, \ldots, o_i\}$ from the policy. Each output is then evaluated by a customised domain-specific reward functions. These include checks for chemical validity (`smiles_acc`), structural accuracy (`atom_acc`, `func_group_acc`), format compliance (`output_format`), and factual correctness (`answer_acc`). Specifically, the `answer_acc` function is designed to handle the variance answer from CheMM-Bench. For textual answers, such as SMILES and IUPAC names, we calculate a fuzzy similarity score based on the Levenshtein distance, which measures the number of edits (insertions, deletions, or substitutions) required to match two strings. This approach allows the function to reward answers with minor, syntactically trivial errors, such as a misplaced hyphen or a single-character typo, rather than marking them as entirely incorrect. For numerical values like molecular weights, answers are rewarded if they fall within a ± 0.05 g/mol tolerance threshold. This threshold represents a $< 0.02\%$ relative error for typical organic molecules, which is consistent with analytical chemistry precision standards. The final advantage A^i is derived by first calculating a single, comprehensive reward r_i for each output through a weighted sum of all reward function scores. The

accumulated rewards for the group form a vector $\mathbf{r} = \{w_1 r_1, w_2 r_2, \ldots, w_i r_i\}$, which is then normalised to calculate the final advantage:

$$A^i = \frac{r_i - \text{mean}(\{w_1 r_1, w_2 r_2, \ldots, w_i r_i\})}{\text{std}(\{w_1 r_1, w_2 r_2, \ldots, w_i r_i\})} \quad (1)$$

The policy is then optimised by maximising the GRPO objective function, steering the model towards generating responses that are chemically sound, factually accurate, and correctly formatted. **See Appendix C for detailed and completed algorithm and mathematic implementation.**

4.2 Implementation Details

The proposed CheMMGRPO is generally applicable to different MLLMs, to demonstrate the effectiveness of CheMMGRPO, we adopt an open-source and non-reasoning MLLM, Qwen2.5-VL-3B (Bai et al., 2025) and train it with MS-SWIFT (Zhao et al., 2025) framework. For the cold-start phase, we set the maximum model length to $L = 20,480$, the maximum image pixels to $i = 262,144$, and the training batch size to 64. Following the prior study from MS-SWIFT, we use the Adam optimiser with a learning rate of $lr = 1e^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.95$, and $\epsilon = 1e^{-8}$.

For the online policy optimisation phase, we perform 4 rollouts per question with a sampling temperature of 1.0 to encourage diverse reasoning paths. The maximum model length and image pixel count are kept consistent with the cold start phase, while the maximum output sequence length is set to 5,120 to capture the model's complete reasoning paths. Both the reference and policy models are initialised from the cold-started model, with the reference model remaining frozen during RL training. Key hyper-parameters for the policy model includes a learning rate of $1e^{-4}$, a KL divergence coefficient β of 0.001, and a batch size of 48. All experiments are conducted on a cluster of 8 H200 (141GB) GPUs with a 12-hour training time.

The cold-start phase consumes the majority of the CheMM-Bench dataset with 40,000 instances to inject fundamental chemistry knowledge and reasoning patterns into the base model. This 8:1 ratio of SFT data to RL data is aligned with established reasoning model practices (Yao et al., 2024; Guo et al., 2025) and balances comprehensive domain coverage against the heavy computational constraints of multi-rollout GRPO. **See Appendix D for the pipeline visualisation.**

5 Results

We benchmark CheMM-R1 against a series of advanced multimodal large language models, which are categorised as either reasoning or non-reasoning. The reasoning models include proprietary systems such as GPT-O3 (OpenAI, 2025), Gemini-2.5-Pro (Comanici et al., 2025), Claude-4-Sonnet (Anthropic, 2025), Grok-4 (xAI, 2025), and Gemini-2.5-Flash (Comanici et al., 2025), alongside open-source models like GLM-4.1V-Think-9B (Team et al., 2025c) and Kimi-VL-A3B (Team et al., 2025b). For the non-reasoning category, we evaluate the closed-source model GPT-4o and several open-source alternatives: InternVL3-8B (Zhu et al., 2025), InternVL3-38B (Zhu et al., 2025), Gemma3-4B (Team et al., 2025a), Gemma3-27B (Team et al., 2025a), and Qwen2.5-VL-3B (Bai et al., 2025). For a more targeted comparison, we also include ChemVLM-8B (Li et al., 2025), a domain-specific model enhanced with visual molecular recognition. All proprietary models are tested with their default parameters, whereas open-source models are tested with a temperature of 0.0 and a maximum output length of 4,096.

5.1 Main Experimental Results

As shown in Table 2, we evaluate CheMM-R1 against various state-of-the-art (SOTA) MLLMs across all tasks in the CheMM-Bench benchmark, reporting both accuracy and F1 score. To further assess performance, we also conduct a similarity comparison using BLEU and Levenshtein distance for text-base results. As detailed in Table 3, these metrics measure how closely a model's generated answer aligns with the ground truth.

Accuracy results. The average performance for molecule and spectrum of the existing models is all under 15%, and several models such as GLM-4.1V-9B-Thinking, Kimi-VL-A3B, InternVL3-8B, Gemma3-4B, and Qwen-2.5-VL-3B provide no correct answers for the evaluation. Gemini-2.5-Pro shows the best molecule performance across all existing models with an average accuracy of 16.13%, while Grok-4 has the best spectrum performance with an average accuracy of 4.05%. We observe that models tend to guess the final weight score without correctly understanding the molecular structure. To investigate this and ensure a genuine assessment of reasoning, we compare performance using two metrics: one that accepts the final answer only, and a stricter criterion where

Model	SR										SE		Overall Average	
	SmilesQA		IupacQA		MwQA				Average		SpectraQA			
	ACC	F1	ACC	F1	ACC w/o S	ACC w S	F1 w/o S	F1 w S	ACC	F1	ACC	F1	ACC	F1
Reasoning														
GPT-o3	1.62	3.19	0.19	0.39	15.82	0.61	27.32	1.21	5.78	10.94	1.50	2.96	3.34	6.46
Gemini-2.5-Pro	8.37	15.44	3.09	6.00	**37.53**	5.27	**54.57**	10.02	16.13	27.78	1.80	3.54	7.95	14.72
Claude-Sonnet-4	1.42	2.80	2.32	4.53	2.23	0.20	4.37	0.40	1.99	3.91	1.60	3.15	1.77	3.48
Grok-4	1.01	2.01	0.77	1.53	6.69	0.41	12.55	0.81	2.79	5.43	4.05	7.78	3.51	6.78
Gemini-2.5-Flash	0.61	1.21	3.09	5.99	22.31	0.81	36.48	1.61	8.58	15.80	1.10	2.18	4.31	8.26
GLM-4.1V-9B-Thinking	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Kimi-VL-A3B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Non-Reasoning														
GPT-4o	0.41	0.81	0.39	0.77	1.01	0.00	2.01	0.00	0.60	1.19	0.10	0.20	0.31	0.63
InternVL3-8B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
InternVL3-38B	0.00	0.00	0.00	0.00	0.41	0.00	0.81	0.00	0.13	0.27	0.05	0.11	0.09	0.18
Gemma3-4B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Gemma3-27B	0.00	0.00	0.77	1.53	0.20	0.00	0.40	0.00	0.33	0.66	0.05	0.10	0.17	0.35
ChemVLM-8B	1.01	2.01	0.00	0.00	0.00	0.00	0.00	0.00	0.33	0.66	-	-	-	-
Qwen-2.5-VL-3B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CheMM-R1 (Ours)	**43.51**	**60.63**	**6.80**	**12.73**	21.95	**15.45**	36.00	**26.76**	**23.73**	**38.35**	**36.32**	**53.28**	**30.92**	**47.23**

Table 2: Main experimental results. Experiments are conducted with CheMM-R1 (3B) on all tasks on CheMM-Bench, and compare with various SOTA reasoning and non-reasoning MLLMs. "SR" is Structure Recognition and "SE" is Structure Elucidation. In MwQA, "w/o S" means without SMILES and "w S" means with SMILES.

both the generated SMILES structure and the final answer must be accurate. This stricter criterion reveals a dramatic drop in performance. For example, the accuracy of Gemini-2.5-pro plummets from 37.53% to 5.27%, while that of Gemini-2.5-Flash drops from 22.31% to just 0.81%. This indicates that existing MLLMs lack the multimodal chemical knowledge for providing accurate compound information. Proprietary models generally perform better than open-source models in both with and without reasoning settings. Gemma3-27B is the best-performing open-source model with an overall accuracy of 0.17%.

Our CheMM-R1 outperforms all existing models on both structure recognition and structure elucidation tasks, achieving an average accuracy of 30.92%. Comparing with Gemini-2.5-Pro, CheMM-R1's performance is only worse on the MwQA task when the evaluation considers just the final answer. However, when applying a stricter criterion that requires both the molecular structure and the final answer to be correct, CheMM-R1's accuracy of 15.45% significantly surpasses Gemini-2.5-Pro's 5.27%. We hypothesise that current reasoning and non-reasoning MLLMs lack the specialised chemistry knowledge to recognise chemical structures from 2D annotations or to elucidate compounds from experimental spectra.

Similarity results. To further evaluate performance on textual chemical questions, we conduct a similarity analysis for the structure recognition and elucidation tasks. We calculate both the BLEU (Papineni et al., 2002) score with uni-gram and the Levenshtein distance (Levenshtein, 1966) with 0.9 threshold (LD-0.9) between the model-generated answers and the ground truth. As shown in Table 3, the trends in similarity largely mirror the accuracy results. Gemini-2.5-Pro outperforms other existing models on both the SmilesQA and SpectraQA tasks by a notable margin. Among the open-source models, Gemma3-27B ranks as the top performer. However, an interesting exception occurs on the IupacQA task, where the answer similarity of Gemini-2.5-Flash surpasses that of Gemini-2.5-Pro on both metrics. This suggests that while Gemini-2.5-Flash can generate answers that are textually close to the correct output, it struggles to translate this into a correct understanding of the underlying chemical

Model	SR				SE	
	SmilesQA		IupacQA		SpectraQA	
	BLEU-1	LD-0.9	BLEU-1	LD-0.9	BLEU-1	LD-0.9
Reasoning						
GPT-o3	0.27	4.26	1.86	9.65	0.59	3.90
Gemini-2.5-Pro	8.37	23.67	13.39	15.47	1.80	5.20
Claude-Sonnet-4	0.20	5.48	9.09	10.42	0.04	6.35
Grok-4	0.25	4.87	7.87	5.60	0.09	7.50
Gemini-2.5-Flash	0.18	6.09	15.03	21.04	0.04	3.15
GLM-4.1V-9B-Thinking	0.00	0.00	0.14	0.00	0.00	0.00
Kimi-VL-A3B	0.00	0.00	0.06	0.00	0.00	0.00
Non-Reasoning						
GPT-4o	0.10	1.83	8.37	1.54	0.02	0.15
InternVL3-8B	0.00	0.00	2.24	0.00	0.00	0.00
InternVL3-38B	0.00	0.00	3.16	0.00	0.00	0.11
Gemma3-4B	0.00	0.00	4.42	0.55	0.00	0.00
Gemma3-27B	0.00	1.42	10.62	3.67	0.05	0.25
ChemVLM-8B	1.01	1.83	0.00	0.00	-	-
Qwen-2.5-VL-3B	0.00	0.61	6.47	0.00	0.00	0.00
CheMM-R1 (Ours)	**43.06**	**68.87**	**28.37**	**23.30**	**36.14**	**63.23**

Table 3: Similarity results. Similarity comparison is conducted with CheMM-R1 (3B) and various SOTA reasoning and non-reasoning MLLMs on all four tasks on CheMM-Bench. "BLEU-1" is uni-gram BLEU score. "LD-0.9" is Levenshtein distance with 90% threshold.

Method	ACC	F1	BLEU	LD-0.9
Cold Start	28.72	44.62	28.60	49.90
Cold Start + GRPO	39.18	56.30	39.15	64.29
GRPO	1.23	2.43	1.23	14.37
CheMMGRPO	0.00	0.00	0.00	3.05
Cold Start + CheMMGRPO (Ours)	**43.51**	**60.63**	**43.06**	**68.87**

Table 4: Ablation study results with SmilesQA. Different settings of model training has been implemented and compared with CheMM-R1.

bonding and structure.

In terms of similarity performance, our CheMM-R1 outperforms all existing models by a notable margin. The performance gap is particularly evident when compared to other leading models. For instance, against Gemini-2.5-Flash, CheMM-R1 achieves scores of 28.37% on BLEU-1 and 23.30% on Levenshtein distance, surpassing Gemini-2.5-Flash's scores of 15.03% and 21.04%, respectively. Although CheMM-R1 achieves state-of-the-art performance among existing MLLMs, its overall accuracy remains far from perfect. This performance gap highlights the challenge that persists in enabling models to fully comprehend chemistry.

Tanimoto similarity While SMILES matching and its similarity are our primary criterion to evaluate precise reasoning, it can penalise models that generate chemically valid results with sequence variations. To provide a unbiased view, we also evaluate the models using Tanimoto similarity at a 1.0 threshold. Under this metric, models demonstrate a stronger grasp of structural concepts. As shown in Table 5, in the SmilesQA task, Gemini-2.5-Pro achieves a 44.52% Tanimoto score, and our CheMM-R1 achieves 60.00%. For SpectraQA, CheMM-R1 achieves a 56.57% Tanimoto score (Feldmann and Bajorath, 2022). This confirms that while current models struggle with the exact precision required for chemistry tasks, CheMM-R1 successfully learns and generates highly meaningful chemical structures.

5.2 Ablation Study

We conduct an ablation study of CheMMGRPO on the SmilesQA dataset to examine the effectiveness of our multimodal chemical reasoning rewards and the role of the cold start phase. As shown in Table 4, our complete method, combining cold start with CheMMGRPO, achieves the best performance across all metrics. The ablation results first underscore the necessity of the cold start phase for imparting foundational knowledge. Without this initial knowledge injection, both standard GRPO and CheMMGRPO completely fail, achieving near-zero accuracy and F1 scores, since they are unable to comprehend complex chemical compounds or produce valid SMILES. Furthermore, the results highlight the value of our domain-specific rewards. When initialise with a cold start, CheMMGRPO outperforms a standard GRPO model by 4.33% in accuracy and 15.92% in LD-0.9 score. Our complete method (CheMM-R1) demonstrates a clear synergistic effect, more than doubling the accuracy and achieving a 18.97% relative improvement in LD-0.9 score comparing with the cold start alone. This demonstrates that CheMMGRPO's approach of reinforcing chemical validity, structural accuracy, format compliance, and factual correctness, combined with cold start, effectively improves for complex multimodal chemistry reasoning tasks.

Model	SmilesQA Tanimoto@1.0	SpectraQA Tanimoto@1.0
Reasoning		
GPT-o3	17.50	7.85
Gemini-2.5-Pro	44.52	11.46
Claude-Sonnet-4	9.42	3.60
Grok-4	2.18	18.28
Gemini-2.5-flash	27.07	5.53
GLM-4.1V-9B-Thinking	0.00	0.00
Kimi-VL-A3B	0.00	0.00
Non-Reasoning		
GPT-4o	2.54	0.26
InternVL3-8B	0.00	0.00
InternVL3-38B	0.00	0.07
Gemma3-4B	0.00	0.00
Gemma3-27B	0.00	0.07
ChemVLM-8B	31.99	-
Qwen-2.5-VL-3B	0.00	0.00
CheMM-R1 (Ours)	**60.00**	**56.57**

Table 5: Tanimoto similarity comparison @1.0 threshold, across reasoning and non-reasoning MLLMs.

6 Conclusion

We introduce CheMM-Bench, a multimodal reasoning benchmark for evaluating MLLMs on chemistry structure recognition and elucidation. Unlike existing datasets limited to single molecular images, CheMM-Bench incorporates multiple spectral images, enabling broader chemistry knowledge assessment. We further introduce CheMMGRPO, a domain-specific reinforcement learning method that guides models toward chemically valid and accurate responses and CheMM-R1, a MLLM trained with CheMMGRPO, establishes new SOTA performance on all CheMM-Bench tasks.

Limitations

In this study, due to computational constraints, CheMM-R1 is limited to 3 billion parameters, which hindered its performance on complex reasoning tasks. Similarly, we are unable to adopt a Mixture of Experts model, imposing an additional limitation on the model's reasoning capacity. We encourage future research to employ larger MLLMs or incorporate Mixture of Experts architectures to further strengthen reasoning capabilities.

Ethical Considerations

All experiments strictly adhere to the Code of Ethics. The original data are public available from previous research (Alberts et al., 2024b). We ensure that our data collection and analysis processes align with ethical guidelines and data copyright regulations.

Acknowledgements

This research was supported by the use of the Katana high-performance computing cluster at University of New South Wales Sydney (Katana, 2010).

References

Marvin Alberts, Teodoro Laino, and Alain C. Vaucher. 2024a. Leveraging infrared spectroscopy for automated structure elucidation. *Communications Chemistry*, 7(1):268.

Marvin Alberts, Oliver Schilter, Federico Zipoli, Nina Hartrampf, and Teodoro Laino. 2024b. Unraveling molecular structure: A multimodal spectroscopic dataset for chemistry. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.

Anthropic. 2025. Introducing claude 4.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. Qwen2.5-vl technical report. *Preprint*, arXiv:2502.13923.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3290 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *Preprint*, arXiv:2507.06261.

Henri A. Favre and Warren H. Powell. 2013. *Nomenclature of Organic Chemistry*. The Royal Society of Chemistry.

Christian Feldmann and Jürgen Bajorath. 2022. Calculation of exact Shapley values for support vector machines with Tanimoto kernel enables model interpretation. *iScience*, 25(9):105023.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.

Kehan Guo, Bozhao Nan, Yujun Zhou, Taicheng Guo, Zhichun Guo, Mihir Surve, Zhenwen Liang, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Can llms solve molecule puzzles? a multimodal benchmark for molecular structure elucidation. In *Advances in Neural Information Processing Systems*, volume 37, pages 134721–134746. Curran Associates, Inc.

Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2025. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *CoRR*, abs/2503.06749.

Zhaorui Huang, Michael S. Chen, Cristian P. Woroch, Thomas E. Markland, and Matthew W. Kanan. 2021. A framework for automated structure elucidation from routine nmr spectra. *Chem. Sci.*, 12:15329–15338.

Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 79 others. 2024. Gpt-4o system card. *CoRR*, abs/2410.21276.

Guwon Jung, Son Gyo Jung, and Jacqueline M. Cole. 2023. Automatic materials characterization from infrared spectra using convolutional neural networks. *Chem. Sci.*, 14:3600–3609.

UNSW Katana. 2010. Katana.

Sunghwan Kim, Jian Chen, Tiejun Cheng, Asta Gindulyte, Jie He, Shu He, Qian Li, Bradford A. Shoemaker, Paul A. Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan E. Bolton. 2025. Pubchem 2025 update. *Nucleic Acids Research*, 53(D1):D1516–D1525.

David R. Klein. 2013. *Organic Chemistry*. Wiley.

Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710.

Junxian Li, Di Zhang, Xunzhi Wang, Zeying Hao, Jingdi Lei, Qian Tan, Cai Zhou, Wei Liu, Yaotian Yang, Xinrui Xiong, Weiyun Wang, Zhe Chen, Wenhai Wang, Wei Li, Mao Su, Shufei Zhang, Wanli Ouyang, Yuqiang Li, and Dongzhan Zhou. 2025. Chemvlm: Exploring the power of multimodal large language models in chemistry area. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, pages 415–423. AAAI Press.

Daniel M. Lowe. 2012. *Extraction of chemical structures and reactions from the literature*. Ph.D. thesis, University of Cambridge, UK.

Nick Mecklenburg, Yiyou Lin, Xiaoxiao Li, Daniel Holstein, Leonardo Nunes, Sara Malvar, Bruno Silva, Ranveer Chandra, Vijay Aski, Pavan Kumar Reddy Yannam, Tolga Aktas, and Todd Hendry. 2024. Injecting new knowledge into large language models via supervised fine-tuning. *Preprint*, arXiv:2404.00213.

Daylight online. 2024a. Daylight chemical information systems. `https://www.daylight.com/dayhtml/doc/theory/theory.smarts.html`.

RDKit online. 2024b. Rdkit: Open-source cheminformatics. `http://www.rdkit.org`.

OpenAI. 2025. Introducing openai o3 and o4-mini.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Yujie Qian, Jiang Guo, Zhengkai Tu, Zhening Li, Connor W. Coley, and Regina Barzilay. 2023. Molscribe: Robust molecular structure recognition with image-to-graph generation. *J. Chem. Inf. Model.*, 63(7):1925–1934.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025a. Gemma 3 technical report. *Preprint*, arXiv:2503.19786.

Kimi Team, Angang Du, Bohong Yin, Bowei Xing, Bowen Qu, Bowen Wang, Cheng Chen, Chenlin Zhang, Chenzhuang Du, Chu Wei, Congcong Wang, Dehao Zhang, Dikang Du, Dongliang Wang, Enming Yuan, Enzhe Lu, Fang Li, Flood Sung, Guangda Wei, and 76 others. 2025b. Kimi-vl technical report. *Preprint*, arXiv:2504.07491.

V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihan Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, and 69 others. 2025c. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *Preprint*, arXiv:2507.01006.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA. Curran Associates Inc.

Lai Wei, Yuting Li, Kaipeng Zheng, Chen Wang, Yue Wang, Linghe Kong, Lichao Sun, and Weiran Huang. 2025. Advancing multimodal reasoning via reinforcement learning with cold start. *Preprint*, arXiv:2505.22334.

David Weininger. 1988. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.*, 28(1):31–36.

Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun, Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, and 8 others. 2024. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *CoRR*, abs/2412.10302.

xAI. 2025. Grok 4.

Yuxi Xie, Anirudh Goyal, Wenyue Zheng, Min-Yen Kan, Timothy P. Lillicrap, Kenji Kawaguchi, and Michael Shieh. 2024. Monte carlo tree search boosts reasoning via iterative preference learning. *CoRR*, abs/2405.00451.

Huanjin Yao, Jiaxing Huang, Wenhao Wu, Jingyi Zhang, Yibo Wang, Shunyu Liu, Yingjie Wang, Yuxin Song, Haocheng Feng, Li Shen, and Dacheng Tao. 2024. Mulberry: Empowering MLLM with o1-like reasoning and reflection via collective monte carlo tree search. *CoRR*, abs/2412.18319.

Jingyi Zhang, Jiaxing Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. 2025. R1-VL: learning to reason with multimodal large language models via step-wise group relative policy optimization. *CoRR*, abs/2503.12937.

Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. 2025. SWIFT: A scalable lightweight infrastructure for fine-tuning. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, pages 29733–29735. AAAI Press.

Zihan Zhao, Bo Chen, Jingpiao Li, Lu Chen, Liyang Wen, Pengyu Wang, Zichen Zhu, Danyang Zhang, Yansi Li, Zhongyang Dai, Xin Chen, and Kai Yu. 2024. Chemdfm-x: towards large multimodal model for chemistry. *Sci. China Inf. Sci.*, 67(12).

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *Preprint*, arXiv:2504.10479.

Task		Text Instance		Image
		Number	Avg Length	Instance
SR	SmilesQA	6,000	3,739.49	6,000
	IupacQA	10,250	5,095.34	10,250
	MwQA	10,250	3,983.93	10,250
SE	SpectraQA	22,000	13,441.04	110,000

Table 6: The number of text and image instances and the average length of the text instances for each dataset.

A Overall Data Statistics

Table 6 provides an overview of the text and image instance counts for the structure recognition and structure elucidation tasks, along with the corresponding average length of text instances. The SmilesQA dataset for structure recognition comprises 6,000 text instances, whereas the IupacQA and MwQA datasets each consist of 10,250 text instances. The average text length ranges from approximately 3,700 to 5,100 tokens, with image instances aligned in a one-to-one correspondence with the text. In contrast, the structure elucidation task comprises 22,000 text instances, each associated with five image instances, resulting in a total of 110,000 images. The average text length in structure elucidation is significantly longer, at approximately 13,400 tokens, highlighting both the complexity of spectral data and the detailed information required for elucidation.

B Functional Group Analysis

We employ RDKit's (online, 2024b) SMARTS (online, 2024a) pattern matching to analyse the number of functional groups within molecular structures, counting each specific group only once per molecule. The same procedure is used when calculating the functional group reward in CheMMGRPO, where functional groups is extracted directly from the molecular structures. Table 7 lists the SMARTS patterns employed in this work, which are derived from previous research (Jung et al., 2023).

C CheMMGRPO Algorithm

We present a CheMMGRPO algorithm, as shown in Algorithm 1, tailored for chemistry-related tasks, particularly IUPAC naming, SMILES generation, and molecular property prediction. The algorithm operates by sampling multiple responses from the current policy for each input query, creating a group of candidate solutions with configurable group size G. Each response is evaluated using five distinct reward functions with specific scoring mechanisms. The answer accuracy function, weighted at 4.0, employs task-specific evaluation: for IUPAC and SMILES tasks, exact matches receive a score of 1.0 while partial matches use fuzzy string matching with Levenshtein Distance to assign scores between 0 and 1 based on string similarity; for molecular weight predictions, answers within ±0.05 of the ground truth receive 1.0 and others receive 0.0. The SMILES validity function, weighted at 1.0, validates the syntactic correctness of generated SMILES strings using RDKit's (online, 2024b) molecular parser, assigning 1.0 for valid structures with at least one atom and 0.0 otherwise. The atom count matching function, weighted at 1.0, parses both student and ground truth SMILES strings to count specific atoms including carbon, oxygen, nitrogen, sulfur, phosphorus, silicon, boron, and halogens, awarding 1.0 for exact count matching across all atom types and 0.0 otherwise. The functional group correctness function, weighted at 1.0, identifies functional groups in both student and reference SMILES using pattern matching, then computes the percentage of correctly identified groups as the ratio of intersection size to ground truth group count. The format compliance function, weighted at 3.0, validates that responses follow the required structure with `<think>`, `<smiles>`, and `<answer>` tags in the correct order, assigning 1.0 for proper format and 0.0 otherwise. The weighted rewards are combined and adjusted with a KL divergence penalty term to prevent the policy from deviating too far from a reference model. Unlike traditional reinforcement learning methods that rely on value networks, CheMMGRPO computes advantages by comparing each response's reward against the group's mean reward, normalised by the group's standard deviation. This group-relative advantage estimation eliminates the need for a separate critic network while providing stable gradient estimates. The policy is then updated using policy gradient methods, where responses with above-average rewards are reinforced and below-average responses are penalised. This approach is particularly effective for chemistry tasks where multiple valid solutions may exist and relative quality assessment within a group provides more robust learning signals than absolute reward values.

Functional Group	SMARTS Pattern
Acid anhydride	[CX3](=[OX1])[OX2][CX3](=[OX1])
Acyl halide	[CX3](=[OX1])[F,Cl,Br,I]
Alcohol	[#6][OX2H]
Aldehyde	[CX3H1](=O)[#6,H]
Alkane	[CX4;H3,H2]
Alkene	[CX3]=[CX3]
Alkyne	[CX2]#[CX2]
Amide	[NX3][CX3](=[OX1])[#6]
Amine	[NX3;H2,H1,H0;!$(NC=O)]
Arene	[cX3]1[cX3][cX3][cX3][cX3][cX3]1
Azo compound	[#6][NX2]=[NX2][#6]
Carbamate	[NX3][CX3](=[OX1])[OX2H0]
Carboxylic acid	[CX3](=O)[OX2H]
Enamine	[NX3][CX3]=[CX3]
Enol	[OX2H][#6X3]=[#6]
Ester	[#6][CX3](=O)[OX2H0][#6]
Ether	[OD2]([#6])[#6]
Haloalkane	[#6][F,Cl,Br,I]
Hydrazine	[NX3][NX3]
Hydrazone	[NX3][NX2]=[#6]
Imide	[CX3](=[OX1])[NX3][CX3](=[OX1])
Imine	$(([CX3]([#6])[#6]),$([CX3H][#6]))]=[$(([NX2][#6]),$([NX2H]))]
Isocyanate	[NX2]=[C]=[O]
Isothiocyanate	[NX2]=[C]=[S]
Ketone	[#6][CX3](=O)[#6]
Nitrile	[NX1]#[CX2]
Phenol	[OX2H][cX3]:[c]
Phosphine	[PX3]
Sulfide	[#16X2H0]
Sulfonamide	[#16X4]([NX3])(=[OX1])(=[OX1])[#6]
Sulfonate	[#16X4](=[OX1])(=[OX1])([#6])[OX2H0]
Sulfone	[#16X4](=[OX1])(=[OX1])([#6])[#6]
Sulfonic acid	[#16X4](=[OX1])(=[OX1])([#6])[OX2H]
Sulfoxide	[#16X3]=[OX1]
Thial	[CX3H1](=S)[#6,H]
Thioamide	[NX3][CX3]=[SX1]
Thiol	[#16X2H]

Table 7: SMARTS pattern for functional group analysis.

Algorithm 1 CheMM GRPO (Group Relative Policy Optimisation)

Require: Policy model π_θ, reference model π_{ref}, dataset $\mathcal{D}$
Require: Reward weights: $w_{\text{acc}} = 4.0$, $w_{\text{smiles}} = 1.0$, $w_{\text{atom}} = 1.0$, $w_{\text{func}} = 1.0$, $w_{\text{fmt}} = 3.0$
Require: Group size G, learning rate α, KL penalty coefficient β
Ensure: Optimized policy π_θ

Initialize policy parameters θ from pre-trained model
for each training iteration **do**
  for each query $q \in \mathcal{D}$ **do**
    // Sample multiple responses per query
    for $i = 1$ to G **do**
      Sample response $y_i \sim \pi_\theta(\cdot|q)$
      Compute log probability $\log \pi_\theta(y_i|q)$
      Compute reference log probability $\log \pi_{\text{ref}}(y_i|q)$
    end for
    // Calculate individual reward components
    for $i = 1$ to G **do**
      $r^{(i)}_{\text{acc}} \leftarrow \text{AnswerScore}(y_i, \text{ground truth})$ ▷ Exact/fuzzy match
      $r^{(i)}_{\text{smiles}} \leftarrow \text{SmilesScore}(y_i)$ ▷ SMILES validity
      $r^{(i)}_{\text{atom}} \leftarrow \text{AtomCountScore}(y_i, \text{ground truth})$ ▷ Atom count match
      $r^{(i)}_{\text{func}} \leftarrow \text{FuncGroupScore}(y_i, \text{ground truth})$ ▷ Functional groups
      $r^{(i)}_{\text{fmt}} \leftarrow \text{FormatScore}(y_i)$ ▷ Format correctness
    end for
    // Compute weighted total rewards
    for $i = 1$ to G **do**
      $r^{(i)} \leftarrow w_{\text{acc}} \cdot r^{(i)}_{\text{acc}} + w_{\text{smiles}} \cdot r^{(i)}_{\text{smiles}} + w_{\text{atom}} \cdot r^{(i)}_{\text{atom}}$
        $+ w_{\text{func}} \cdot r^{(i)}_{\text{func}} + w_{\text{fmt}} \cdot r^{(i)}_{\text{fmt}}$
      $\text{KL}^{(i)} \leftarrow \log \pi_\theta(y_i|q) - \log \pi_{\text{ref}}(y_i|q)$ ▷ KL divergence
      $r^{(i)}_{\text{total}} \leftarrow r^{(i)} - \beta \cdot \text{KL}^{(i)}$ ▷ KL-penalised reward
    end for
    // Calculate group statistics
    $\bar{r} \leftarrow \frac{1}{G} \sum_{i=1}^{G} r^{(i)}_{\text{total}}$ ▷ Mean reward
    $\sigma_r \leftarrow \sqrt{\frac{1}{G} \sum_{i=1}^{G} (r^{(i)}_{\text{total}} - \bar{r})^2}$ ▷ Std deviation
    // Compute normalised advantages
    for $i = 1$ to G **do**
      $A^{(i)} \leftarrow \frac{r^{(i)}_{\text{total}} - \bar{r}}{\sigma_r + \epsilon}$ ▷ Normalised advantage
    end for
    // Compute policy gradient
    $\nabla_\theta J \leftarrow \frac{1}{G} \sum_{i=1}^{G} A^{(i)} \cdot \nabla_\theta \log \pi_\theta(y_i|q)$
    // Update policy parameters
    $\theta \leftarrow \theta + \alpha \cdot \nabla_\theta J$
  end for
end for
return Optimised policy π_θ

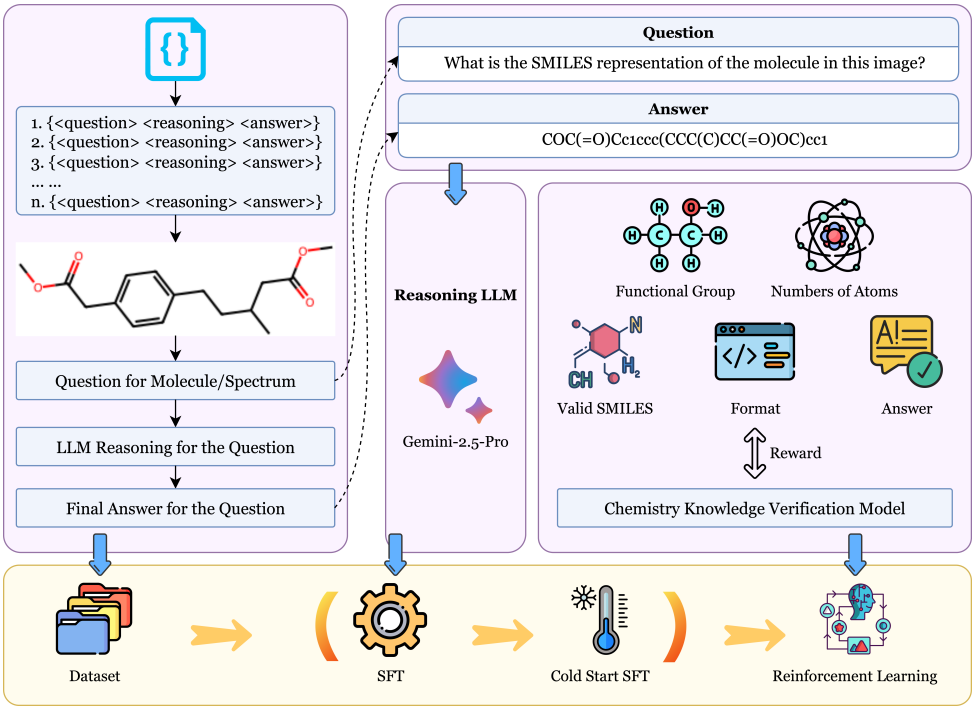

Figure 4: Pipeline overview of training CheMM-R1.

D CheMM-R1 Pipeline Overview

We present the training pipeline of CheMM-R1 in Figure 4. The pipeline begins with the collection of CheMM-Bench, where we extract reasoning paths from Gemini-2.5-Pro using Chain-of-Thought prompts with seed question-answer pairs. The collected data is used for cold-start supervised fine-tuning based on the Qwen-2.5-VL-3B model. During cold start, the model learns fundamental chemistry knowledge and critical thinking steps leading to the final answer. After supervised fine-tuning, the model is trained with CheMMGRPO, which employs five reward functions to assess chemical validity, structural accuracy, format compliance, and factual correctness. CheMMGRPO further improves CheMM-R1's performance by optimizing the model toward generating chemically sound and accurate responses.

E Prompts

We present example prompts for all tasks in CheMM-Bench in Figures 5–8. Each example instance contains three main sections. The first section comprises the original question-image pairs, including molecular structure images for structure recognition tasks and spectral images for structure elucidation tasks. The second section contains task-specific Chain-of-Thought prompts designed to guide reasoning path generation. These prompts, along with the question, image, and answer, are sent to Gemini-2.5-Pro for reasoning path generation. The third section presents the generated reasoning steps and final answer from Gemini-2.5-Pro. All generations are saved in a `JSON` file, with corresponding SMILES annotations formatted using `<think>`, `<smiles>`, and `<answer>` tags.

SMILES

Question:

What is the SMILES representation of the molecule in this image?

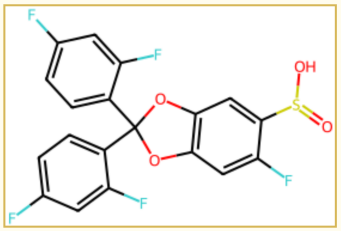

Reasoning Steps:

1. Identify all atoms present in the molecule and their exact connectivity.
2. Determine the presence of any functional groups, rings, or specific structural patterns.
3. Note the stereochemistry, charge states, and isotopic information where visible.
4. Pay special attention to bond types (single, double, triple, aromatic)
5. Explain how each part of the molecule is represented in the SMILES notation.
6. Clarify any complex notations or conventions used in the SMILES string.

Reasoning Outputs:

1. Deconstruction of the Molecular Structure: The molecule consists of a central fused ring system, which is a dihydrobenzo[b][1,4]dioxine. This core structure is heavily substituted. The benzene part of the core has a sulfonic acid (-SO3H) group and a fluorine (F) atom. The dihydrodioxine part contains a spiro carbon atom, which is bonded to two identical 2,4-difluorophenyl groups.
2. Representation of the Sulfonic Acid Group: The SMILES string begins with 'O=S(O)'. This represents the sulfonic acid group attached to the benzene ring. 'S' is the sulfur atom, 'O=' indicates a double bond to one oxygen, and '(O)' indicates a single bond to a hydroxyl group. This is a compact representation of the -SO3H functional group.
3. Representation of the Benzene Portion of the Fused System: The benzene ring with its substituents and fusion points is represented by the sequence 'c1cc2c(cc1F)'.
- 'c1' signifies the first atom of an aromatic ring, which is the carbon atom bonded to the sulfonic acid group 'O=S(O)'.
- The traversal continues along the ring through 'c' and then to 'c2'. 'c2' is an atom in the benzene ring that is also a fusion point for the second ring (the dihydrodioxine ring). The number '2' is a label for this fusion point.
- The next part, 'c(cc1F)', represents the second fusion point and the remaining path of the benzene ring. This complex notation encodes the connectivity and substitution, indicating that the fluorine atom ('F') is located meta to the sulfonic acid group, which matches the visual structure.
4. Representation of the Dihydrodioxine Ring: The atoms of the second ring are described between the two fusion points. In this SMILES, the path for the second ring starts after 'c2'. The description is 'OC(c1ccc(F)cc1F)(c1ccc(F)cc1F)O2'.
- The path starts from the second fusion atom (encoded in 'c(cc1F)') which is bonded to an oxygen ('O').
- This oxygen is bonded to the spiro carbon 'C'.
- This spiro carbon has two branches, each represented by '(c1ccc(F)cc1F)'.
- The spiro carbon is then bonded to a second oxygen atom 'O'.
- This final oxygen atom has the ring-closing label '2' ('O2'), indicating it forms a bond back to the atom labeled 'c2', thus completing the six-membered dihydrodioxine ring.
5. Representation of the 2,4-Difluorophenyl Groups: The two identical branches on the spiro carbon are represented by '(c1ccc(F)cc1F)'. This is the SMILES for a 2,4-difluorophenyl group. While there are multiple ways to write this, the given string is a valid, albeit non-canonical, representation. It can be interpreted as a traversal around the phenyl ring starting from the attachment point ('c1'), proceeding in a way ('ccc(F)cc1F') that places the fluorine atoms at positions 2 and 4 relative to the attachment point, as shown in the image.
6. Overall Assembly: By combining these fragments, the SMILES string 'O=S(O)c1cc2c(cc1F)OC(c1ccc(F)cc1F)(c1ccc(F)cc1F)O2' provides a complete linear representation of the complex, multi-ring, and substituted molecule.

Final Answer:

O=S(O)c1cc2c(cc1F)OC(c1ccc(F)cc1F)(c1ccc(F)cc1F)O2

Figure 5: Chain-of-Though prompt and reasoning path of SmilesQA.

IUPAC Name

Question:

What is the IUPAC name of the molecule shown in the image?

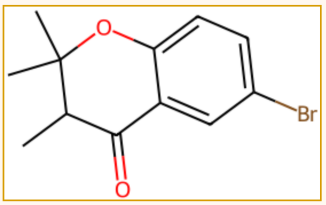

Reasoning Steps:

Part 1:From Visual Structure to SMILES
1. Identify Key Features: From the image, identify all atoms, their connectivity, functional groups, rings, and any specific structural patterns.
2. Note Stereochemistry and Bonds: Note any stereochemistry (e.g., wedges, dashes), charge states, isotopic information, and bond types (single, double, triple, aromatic).
3. Explain the SMILES String: Explain how each part of the molecule (atoms, bonds, branches, rings) is represented in the provided SMILES notation. Clarify any complex notations or conventions used.

Part 2: From SMILES to IUPAC Name
4. Identify the Principal Functional Group: Based on the SMILES string, determine the principal functional group. This will define the suffix of the IUPAC name.
5. Find the Parent Chain/Ring: Identify the longest continuous carbon chain (or the main ring system) that contains the principal functional group.
6. Number the Parent Chain: Number the atoms of the parent chain, starting from the end that gives the lowest possible locant (number) to the principal functional group. If applicable, then assign the lowest locants to multiple bonds, and finally to substituents.
7. Identify and Name Substituents: List all the substituent groups attached to the parent chain and name them according to IUPAC rules.
8. Determine Stereochemistry (if any): Analyze the SMILES notation for stereochemical indicators and assign the appropriate E/Z configuration for double bonds or R/S configuration for chiral centers.
9. Assemble the Full IUPAC Name: Construct the final name by arranging the components in the correct order: stereochemical descriptors, alphabetized substituents with their locants, the parent chain name, and the suffix with its locant.

Reasoning Outputs:

Part 1: From Visual Structure to SMILES
1. Identify Key Features: The image displays a bicyclic molecule composed of two fused six-membered rings. One ring is a benzene ring. The other is a heterocyclic ring containing one oxygen atom. This core structure is a chromene derivative. The molecule has several functional groups and substituents: a ketone group (C=O) on the heterocyclic ring, a bromine atom (Br) on the benzene ring, two methyl groups on one carbon of the heterocyclic ring (a gem-dimethyl group), and one methyl group on an adjacent carbon.
2. Note Stereochemistry and Bonds: The structure is depicted as flat, without any stereochemical indicators like wedges or dashes. The benzene ring contains aromatic bonds. The heterocyclic ring contains single bonds and one carbon-oxygen double bond (ketone). The two rings are fused.
3. Explain the SMILES String: The provided SMILES string is 'CC1C(=O)c2cc(Br)ccc2OC1(C)C'.
- 'CC1...': This starts with a methyl group ('C') attached to a carbon ('C') which is designated as part of ring '1'.
- '...C1C(=O)...': The first ring carbon ('C1') is bonded to another carbon which has a double bond to an oxygen ('=O'), representing the ketone.
- '...C(=O)c2...': The ketone's carbon is bonded to an aromatic carbon ('c') which is part of a second ring, designated as '2'.
- '...c2cc(Br)ccc2...': This sequence describes the benzene ring. It's a six-membered aromatic ring ('c...c') where one carbon has a bromine substituent '(Br)'. The '2' after the last 'c' indicates that this carbon closes ring '2' by bonding back to the first atom of the ring ('c2').
- '...ccc2OC1(C)C': This section connects the two rings and completes the first ring. The closing atom of the benzene ring ('c2') is bonded to an oxygen atom ('O'). This oxygen is then bonded to a carbon atom ('C'), which is designated as part of ring '1' ('C1'), closing the first ring. This same carbon atom ('C1') also has two methyl groups attached, represented by '(C)C'.
Part 2: From SMILES to IUPAC Name
4. Identify the Principal Functional Group: The molecule contains a ketone and an ether within a fused ring system. The base structure is a benzopyranone. Specifically, it is a derivative of chromone (benzopyran-4-one). The ketone ('-one') is the principal functional group and dictates the suffix of the name.
5. Find the Parent Ring: The parent structure is the fused bicyclic system. The standard IUPAC name for this system is chromene. With a ketone at position 4, it is a chromen-4-one. In a standard chromen-4-one, there is a double bond between positions 2 and 3. Here, this bond is saturated. This is indicated in the IUPAC name using an 'indicated hydrogen' prefix, '3H-', which specifies that position 3 is saturated. Thus, the parent name is '3H-chromen-4-one'.
6. Number the Parent Chain: The standard numbering for the chromene ring system starts with the oxygen atom as position 1. The numbering proceeds around the heterocyclic ring first, then the benzene ring.
- O = 1
- Carbon with two methyl groups = 2
- Carbon with one methyl group = 3
- Carbonyl carbon = 4
- The numbering continues around the benzene ring from 5 to 8.
- This numbering places the ketone at position 4, consistent with the parent name.
7. Identify and Name Substituents: Based on the numbering:
- A bromine atom is at position 6. This is named '6-bromo'.
- Two methyl groups are at position 2. This is '2,2-dimethyl'.
- One methyl group is at position 3. This is '3-methyl'.
- Combining all methyl groups gives '2,2,3-trimethyl'.
8. Determine Stereochemistry (if any): The carbon at position 3 is a chiral center. However, neither the drawing nor the SMILES string provides stereochemical information (e.g., R/S configuration), so it is omitted from the name.
9. Assemble the Full IUPAC Name: The substituents are listed alphabetically before the parent name. 'Bromo' comes before 'trimethyl'.
- Substituents with locants: '6-bromo-2,2,3-trimethyl'
- Parent name: '3H-chromen-4-one'
- Combining these parts yields the final IUPAC name: '6-bromo-2,2,3-trimethyl-3H-chromen-4-one'.

Final Answer:

6-bromo-2,2,3-trimethyl-3H-chromen-4-one

Figure 6: Chain-of-Though prompt and reasoning path of IupacQA.

Molecular Weight

Question:

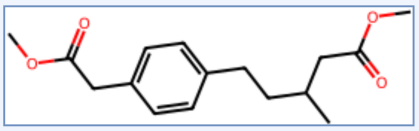

Write the molecular weight of the molecule in the image to 2 decimal places.

Reasoning Steps:

Part 1:From Visual Structure to SMILES
1. Identify Key Features: From the image, identify all atoms, their connectivity, functional groups, rings, and any specific structural patterns.
2. Note Stereochemistry and Bonds: Note any stereochemistry (e.g., wedges, dashes), charge states, isotopic information, and bond types (single, double, triple, aromatic).
3. Explain the SMILES String: Explain how each part of the molecule (atoms, bonds, branches, rings) is represented in the provided SMILES notation. Clarify any complex notations or conventions used.

Part 2: From SMILES to Molecular Weight
4. Derive the Molecular Formula: Parse the SMILES string to determine the count of each unique element. This requires counting all explicitly written atoms and inferring the number of implicit hydrogen atoms needed to satisfy the standard valence of each atom. State the final molecular formula (e.g., $C_4H_{10}O$).
5. List Standard Atomic Weights: For each element present in the molecular formula, list its standard atomic weight from the periodic table (e.g., C ≈ 12.011 u, H ≈ 1.008 u, O ≈ 15.999 u).
6. Calculate the Molecular Weight: Show the final calculation clearly. Multiply the count of each element by its atomic weight and sum the results to find the total molecular weight. State the final answer with the appropriate units (u or g/mol).

Reasoning Outputs:

Part 1: From Visual Structure to SMILES
1. Identify Key Features: The image displays a molecule with a central benzene ring that is para-substituted. One substituent is a methyl acetoxy group ('-CH2-C(=O)O-CH3'). The other substituent is a longer chain containing a methyl branch and ending in a methyl ester: '-CH2-CH2-CH(CH3)-CH2-C(=O)O-CH3'.
2. Note Stereochemistry and Bonds: The structure is a 2D representation. It contains single bonds, double bonds within the aromatic ring and the two carbonyl groups ('C=O'). There is a chiral center on the carbon atom holding the methyl group in the longer chain, but its stereochemistry is not specified.
3. Explain the SMILES String: The SMILES string 'COC(=O)Cc1ccc(CCC(C)CC(=O)OC)cc1' represents the molecule as follows:
- 'COC(=O)C': This segment describes the first methyl ester group. 'CO' is a methyl group attached to an ester oxygen. 'C(=O)' is the carbonyl group, and the final 'C' is the carbon it's attached to.
- 'c1ccc(...)cc1': This denotes the para-substituted benzene ring. 'c1' starts the ring, and the final '1' closes it. The lowercase 'c' signifies aromaticity. The first substituent ('COC(=O)C') is attached to the first carbon of the ring.
- '(...)': The parentheses enclose the second substituent, which is attached at the para position (the fourth carbon in the sequence 'c1ccc...').
- 'CCC(C)CC(=O)OC': This is the second substituent. 'CCC' is a propyl chain. '(C)' indicates a methyl group branching off the last carbon of the propyl chain. 'CC(=O)OC' describes the rest of the chain leading to the second methyl ester group. This encoding accurately reflects the connectivity shown in the image.
Part 2: From SMILES to Molecular Weight
4. Derive the Molecular Formula: By parsing the SMILES string 'COC(=O)Cc1ccc(CCC(C)CC(=O)OC)cc1' and counting each atom, including implicit hydrogens needed to satisfy valency rules (C=4, O=2):
- Carbon (C): 1(methyl) + 1(carbonyl) + 1(methylene) + 6(ring) + 3(chain) + 1(branch) + 1(methylene) + 1(carbonyl) + 1(methyl) = 16 carbons.
- Oxygen (O): 2 in the first ester + 2 in the second ester = 4 oxygens.
- Hydrogen (H): 3(methyl) + 2(methylene) + 4(ring) + 2(chain) + 2(chain) + 1(chain) + 3(branch) + 2(methylene) + 3(methyl) = 22 hydrogens.
- The molecular formula is C\u2081\u2086H\u2082\u2082O\u2084.
5. List Standard Atomic Weights:
- Carbon (C): 12.011 u
- Hydrogen (H): 1.008 u
- Oxygen (O): 15.999 u
6. Calculate the Molecular Weight:
- Weight = (16 \u00d7 12.011 u) + (22 \u00d7 1.008 u) + (4 \u00d7 15.999 u)
- Weight = 192.176 u + 22.176 u + 63.996 u
- Weight = 278.348 u
- Rounding to two decimal places, the molecular weight is 278.35 u.

Final Answer:

278.35

Figure 7: Chain-of-Though prompt and reasoning path of MwQA.

Spectra

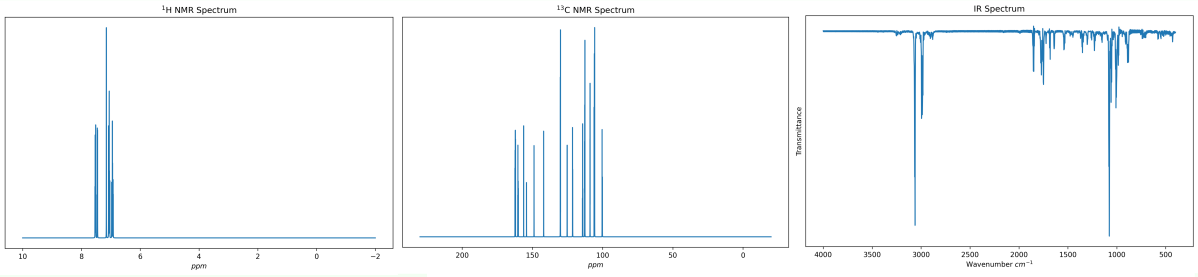

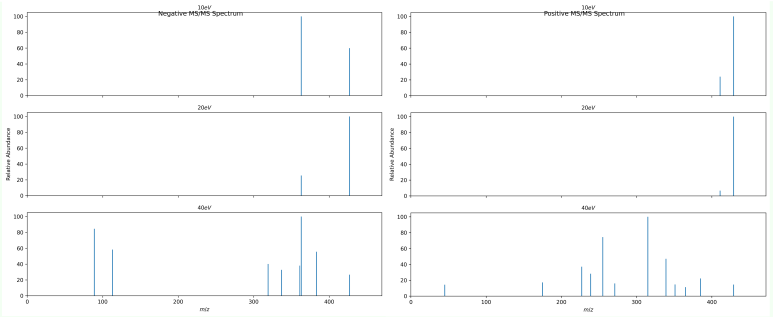

Question:

Given the 1H NMR, 13C NMR, IR, positive-ion mass spectrometry (MS), and negative-ion MS spectra shown in the images, how would you represent this compound in SMILES format?

Reasoning Steps:

Part 1: Mass Spectrometry (MS) Analysis

1. Analyze Molecular Formula: Start with the provided molecular formula (e.g., $C_xH_yN_zO_w$).
A. Calculate Exact Molecular Weight (MW): Calculate the precise molecular weight from the given formula.
B. Calculate Degree of Unsaturation (DoU): Use the formula to calculate the Degree of Unsaturation:
DoU=C - 1/2(H + X - N) +1
State what this number implies about the sum of rings and π-bonds in the structure.

2. Positive-Ion MS Analysis:
A. Identify Molecular Ion: Locate the primary molecular ion, such as the protonated molecule [M+H]+ or a sodium adduct [M+Na]+. Note its mass-to-charge ratio (m/z).
B. Confirm Molecular Weight: Verify that the observed m/z is consistent with the MW calculated from the molecular formula (e.g., show that the m/z of [M+H]+ is approximately MW + 1).
C. Analyze Key Fragments: Identify at least two significant fragment peaks. Explain the likely structural piece that has been lost or that the fragment represents (e.g., a loss of 18 suggests H_2O).

3. Negative-Ion MS Analysis:
A. Identify Molecular Ion: Identify the deprotonated molecule [M−H]−.
B. Confirm Molecular Weight: Show that the m/z of the [M−H]− ion is also consistent with the calculated MW (i.e., m/z = MW - 1).
C. Deduce Functional Group Information: Explain what the presence of a strong [M−H]− signal implies about the molecule's functional groups (e.g., the presence of an acidic proton from a phenol, carboxylic acid, etc.).

Part 2: Infrared (IR) Spectroscopy Analysis
1. Identify Key Absorptions: Pinpoint significant absorption peaks in the IR spectrum by their wavenumber (e.g., 3350 cm^{-1}, 1710 cm^{-1}, 1600 cm^{-1}).
2. Assign Functional Groups: For each key absorption, assign the corresponding functional group (e.g., O-H bond, C=O bond, C=C aromatic bonds). Conclude which functional groups are present in the molecule based on this data.

Part 3: ^{13}C NMR Spectroscopy Analysis
1. Count Carbon Environments: State the total number of signals observed in the ^{13}C NMR spectrum and explain what this indicates about the molecule's symmetry and the number of unique carbon atoms.
2. Analyze Chemical Shifts: For each signal, or for groups of signals in distinct regions (e.g., aromatic, aliphatic), explain what the chemical shift (δ) range tells you about the type of carbon (e.g., sp^3-hybridized C-H, aromatic carbon, carbonyl carbon).

Part 4: ^{1}H NMR Spectroscopy Analysis
1. For each distinct signal in the ^{1}H NMR spectrum, provide a complete analysis:
2. Chemical Shift (δ): Describe what the chemical shift suggests about the proton's electronic environment.
3. Integration: State the relative number of protons the signal represents.
4. Splitting Pattern (Multiplicity): Describe the splitting pattern (e.g., singlet, doublet, triplet) and use the n+1 rule to determine how many neighboring protons are indicated.
5. Assign Fragment: Based on the shift, integration, and splitting, identify the specific structural fragment (e.g., a $-CH_3$ group adjacent to a $-CH_2$ group).

Part 5: Structure Elucidation and SMILES Verification
1. Assemble the Structure: Combine the fragments and clues gathered from all spectra (MS, IR, ^{13}C NMR, and ^{1}H NMR). Explain how these pieces fit together to form a single, coherent structure that is consistent with the molecular formula and the Degree of Unsaturation.
2. Generate and Explain the SMILES String: Based on your deduced structure, generate the corresponding SMILES string. Explain how the final SMILES notation represents the key features of the molecule you have identified (e.g., atoms, branches, rings, bond types).

Reasoning Outputs:

Part 1: Mass Spectrometry (MS) Analysis
A. Positive-Ion Spectrum Analysis
Identify Key Ions: Although a molecular ion like $[M+H]^+$ is not clearly visible in the provided data, we can deduce the molecular weight from the negative-ion spectrum.Determine Molecular Weight (MW): The negative-ion spectrum shows a clear $[M-H]^-$ peak at m/z 173.12955. This indicates the neutral molecule (M) has an integer molecular weight of 174.Propose Molecular Formula: The exact mass of the $[M-H]^-$ peak is 173.12955. Let's test the proposed formula for the given SMILES string, $C_8H_{18}N_2O_2$. The exact mass of $C_8H_{18}N_2O_2$ is 174.1368 u. The expected mass for $[M-H]^-$ ($C_8H_{17}N_2O_2^-$) would be 174.13628 - 1.007825 (H) = 173.129003 u. This is an excellent match for the observed peak at 173.12955, confirming the molecular formula is $C_8H_{18}N_2O_2$.Calculate Degree of Unsaturation (DoU): Using the formula $C_8H_{18}N_2O_2$, the DoU is calculated as: DoU = 8 - (18/2) + (2/2) + 1 = 8 - 9 + 1 + 1 = 1. This indicates the presence of one double bond or one ring in the structure.Analyze Key Fragments:
m/z 57.06988: This peak corresponds to the $C_4H_9^+$ ion, the very stable tert-butyl cation $[(CH_3)_3C]^+$. This is a strong indication of a tert-butoxy or tert-butyl group in the molecule.m/z 73.07602: This peak corresponds to the $C_3H_8N_2^+$ ion, which represents the trimethylhydrazine fragment $[(CH_3)_2N\text{-}N(CH_3)]^+$. This is formed by the cleavage of the N-C bond next to the carbonyl group.

B. Negative-Ion Spectrum Analysis
Identify Key Ion: The most prominent ion at high mass is $[M-H]^-$ at m/z 173.12955. The label suggests a proton is lost from one of the methyl groups.Confirm Molecular Weight: The m/z of 173 confirms that the molecular weight of the neutral compound is 174, consistent with the proposed formula.Deduce Functional Group Information: The ability to form a stable $[M-H]^-$ ion, even by losing a C-H proton, supports the determined molecular weight. A significant fragment is also observed at m/z 73.06589, which corresponds to the tert-butoxide anion $[(CH_3)_3CO]^-$, further confirming the presence of a tert-butoxy group.

Part 2: Infrared (IR) Spectroscopy Analysis
Identify Key Absorptions:
A very strong, sharp absorption is present around 1705 cm^{-1}.Multiple sharp peaks are visible in the 2850-3000 cm^{-1} region.There is an absence of broad peaks in the 3200-3600 cm^{-1} region.Assign Functional Groups:
The intense peak at ~1705 cm^{-1} is characteristic of a C=O (carbonyl) stretch. Its position is consistent with an ester or, more specifically, a carbamate group.The peaks at 2850-3000 cm^{-1} are due to sp^3 C-H stretching from the numerous methyl groups.The absence of O-H or N-H stretching bands confirms there are no alcohol, carboxylic acid, or primary/secondary amine/amide functional groups. This means the two nitrogen atoms must be tertiary.Conclusion: The IR data confirms the presence of a carbonyl group (which accounts for the DoU of 1) and multiple sp^3 C-H bonds, with no acidic O-H or N-H protons.

Part 3: ^{13}C NMR Spectroscopy Analysis
Count Carbon Environments: The spectrum displays 5 distinct signals. Since the molecular formula $C_8H_{18}N_2O_2$ contains 8 carbon atoms, this indicates significant symmetry. The proposed structure, $(CH_3)_2N\text{-}N(CH_3)\text{-}C(=O)O\text{-}C(CH_3)_3$, has exactly 5 unique carbon environments: (a) the two equivalent methyls on one N, (b) the single methyl on the other N, (c) the carbonyl carbon, (d) the quaternary carbon of the t-butyl group, and (e) the three equivalent methyls of the t-butyl group. This matches the data perfectly.Analyze Chemical Shifts:
δ ≈ 154.3 ppm: This downfield signal is characteristic of a carbonyl carbon (C=O) in a carbamate or ester environment.δ ≈ 80.8 ppm: This signal is assigned to the quaternary carbon of the tert-butyl group, which is bonded to the electronegative oxygen atom (-O-$C(CH_3)_3$).δ ≈ 42.5 ppm & δ ≈ 34.1 ppm: These two signals are in the region for sp^3 carbons attached to nitrogen. They correspond to the two different types of N-methyl groups.δ ≈ 28.2 ppm: This upfield signal is characteristic of the three equivalent methyl carbons of the tert-butyl group (-$C(CH_3)_3$).

Part 4: ^{1}H NMR Spectroscopy Analysis
There are three signals, all of which are singlets, indicating no adjacent protons for any of the groups.
Signal at δ ≈ 1.46 ppm:
Chemical Shift: In the aliphatic region, consistent with methyl groups.Integration: 9H, indicating three equivalent methyl groups.Splitting Pattern: A singlet, indicating no adjacent protons.Assign Fragment: This is the classic signature of a tert-butyl group, -$C(CH_3)_3$.Signal at δ ≈ 2.71 ppm:
Chemical Shift: Downfield, suggesting proximity to an electronegative atom like nitrogen.Integration: 6H, indicating two equivalent methyl groups.Splitting Pattern: A singlet.Assign Fragment: This corresponds to the dimethylamino group, $(CH_3)_2N$-.Signal at δ ≈ 2.99 ppm:
Chemical Shift: Furthest downfield methyl signal, suggesting strong electron withdrawal.Integration: 3H, indicating one methyl group.Splitting Pattern: A singlet.Assign Fragment: This is the single methyl group attached to the nitrogen that is also bonded to the carbonyl group, -$N(CH_3)C(=O)$-.

Part 5: Structure Elucidation and SMILES Verification
Assemble the Structure: All the data points to a single structure. We have identified three key fragments from the ^{1}H NMR: a tert-butyl group (-$C(CH_3)_3$), a dimethylamino group ($(CH_3)_2N$-), and a methylamino group (-$N(CH_3)$-). The IR and ^{13}C NMR confirm a carbonyl group, which the DoU requires to be part of an ester/carbamate (-C(=O)O-). Assembling these pieces to satisfy the molecular formula $C_8H_{18}N_2O_2$ and the lack of N-H bonds leads to the structure: $(CH_3)_2N\text{-}N(CH_3)\text{-}C(=O)\text{-}O\text{-}C(CH_3)_3$.Generate and Explain the SMILES String: This structure is represented by the SMILES string CN(C)N(C)C(=O)OC(C)(C)C. Let's break it down:
CN(C): Represents the dimethylamino group, $(CH_3)_2N$-.N(C): Connects the first nitrogen to a second nitrogen, which has a methyl group attached, forming the $(CH_3)_2N\text{-}N(CH_3)$- core.C(=O)O: Connects the second nitrogen to the carbamate group.
• • C(C)(C)C: Represents the tert-butyl group attached to the oxygen. This SMILES string perfectly matches the final deduced structure and the provided answer.

Final Answer:

CN(C)N(C)C(=O)OC(C)(C)C

Figure 8: Chain-of-Though prompt and reasoning path of SpectraQA.