

Boosting LLM’s Molecular Structure Elucidation with Knowledge Enhanced Tree Search Reasoning

Xiang Zhuang^{1,2}, Bin Wu³, Jiyu Cui¹, Kehua Feng^{1,2},
Xiaotong Li^{1,2}, Huabin Xing^{1,2}, Keyan Ding², Qiang Zhang^{1,2*}, Huajun Chen^{1,2*}

¹ Zhejiang University

² ZJU-Hangzhou Global Scientific and Technological Innovation Center

³ University College London

{zhuangxiang, qiang.zhang.cs, huajunsir}@zju.edu.cn

Abstract

Molecular structure elucidation involves deducing a molecule’s structure from various types of spectral data, which is crucial in chemical experimental analysis. While large language models (LLMs) have shown remarkable proficiency in analyzing and reasoning through complex tasks, they still encounter substantial challenges in molecular structure elucidation. We identify that these challenges largely stem from LLMs’ limited grasp of specialized chemical knowledge. In this work, we introduce a Knowledge-enhanced reasoning framework for Molecular Structure Elucidation (K-MSE), leveraging Monte Carlo Tree Search for test-time scaling as a plugin. Specifically, we construct an external molecular substructure knowledge base to extend the LLMs’ coverage of the chemical structure space. Furthermore, we design a specialized molecule-spectrum scorer to act as a reward model for the reasoning process, addressing the issue of inaccurate solution evaluation in LLMs. Experimental results show that our approach significantly boosts performance, particularly gaining more than 20% improvement on both GPT-4o-mini and GPT-4o¹.

1 Introduction

Recent studies have demonstrated the broad application potential of large language models (LLMs) (OpenAI, 2023; Dubey et al., 2024) in chemistry-related tasks, such as reaction prediction (M. Bran et al., 2024) and molecule generation (Liu et al., 2024). However, despite progress in these fields, fully leveraging the deep thinking and reasoning capabilities of LLMs to address complex problems in chemistry, such as molecular structure elucidation, remains a significant challenge. Molecular structure elucidation is a fundamental task in chemical experimental analysis, involving

the deduction of molecular structures from various types of spectral data, such as nuclear magnetic resonance (NMR), infrared (IR) spectroscopy, etc. (Guo et al., 2024). Accurate elucidation of molecular structures is essential for chemical research, as it forms a critical step in interpreting experimental results (Xue et al., 2023). Nevertheless, this task is inherently complex, often requiring significant time and expertise. Even skilled professionals typically need 10 to 15 minutes to reason through the structure of a single molecule (Field et al., 2020). Therefore, automating the interpretation of spectra and accurately deducing molecular structures using LLMs has the potential to greatly enhance experimental efficiency and drive the automation of chemical research (Dai et al., 2024).

Recent studies have demonstrated that increasing the number of reasoning steps, especially with tree search (Hao et al., 2023; Zhang et al., 2024a,b), can greatly improve LLMs’ ability to solve complex problems. This highlights the importance of test-time scaling laws (Snell et al., 2024; OpenAI, 2024c). However, existing methods are not easily applicable to molecular structure elucidation. This is largely due to LLMs’ limited grasp of chemical molecular knowledge (Chen, 2024), which can be attributed to two key factors.

First, LLMs lack comprehensive coverage of chemical molecular structure space. As shown in Figure 1(a), for instance, thiophene—an aromatic heterocycle composed of one sulfur atom and four carbon atoms—is a structure that LLMs often struggle to analyze accurately. Although the model may recognize that the structure contains at most four aromatic carbon atoms and one sulfur atom, it frequently misidentifies it as the benzene ring (the most common aromatic substructure with six carbon atoms), overlooking the limitation on the number of aromatic carbon atoms. This suggests that LLMs lack critical substructure knowledge, hindering their ability to fully comprehend and accurately

*Corresponding author.

¹Code repository: <https://github.com/HICAI-ZJU/K-MSE>.

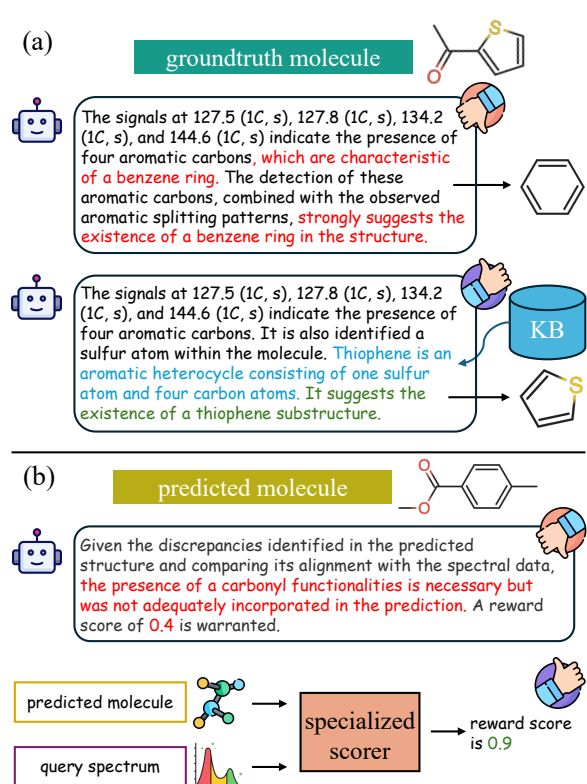


Figure 1: (a) By incorporating a knowledge base, we enhance the ability of LLMs to handle a broader range of chemical molecular structures. (b) Through the development of a specialized scorer, we enable accurate evaluation of potential solutions.

infer complex and unusual chemical structures.

Second, LLMs cannot accurately evaluate and correct their reasoning process. This capability is crucial for tree search-based reasoning, as it guides reasoning to be more effective and efficient (Wan et al., 2024). Timely evaluation and correction provide the necessary feedback, helping LLMs identify potential issues and optimize accordingly. However, we have observed that in molecular structure elucidation tasks, LLMs are significantly deficient in evaluating solutions as a reward model (see Figure 1(b)). This is because reward models require an in-depth understanding of the task domain to reflect the complex relationship between predicted molecules and spectral data. LLMs, lacking domain-specific knowledge, are unable to provide precise guidance in assessment and feedback.

To address these challenges, we propose a Knowledge-enhanced framework based on Monte Carlo Tree Search (MCTS) to improve LLMs’ reasoning for Molecular Structure Elucidation (K-MSE). To tackle the first issue of limited coverage of chemical molecular space in LLMs, we con-

struct an external molecular substructure knowledge base (Figure 1(a)). Recognizing molecular substructures serve as the foundational elements of chemical space, this knowledge base integrates both substructures and their structural descriptions through an automated pipeline. Designed to enhance LLMs with domain-specific chemical knowledge, it improves accuracy in inferring molecular structures and reduces errors in atypical cases.

To overcome the second challenge of LLMs inaccurately evaluating solutions, we design a specialized molecule-spectrum scorer as a reward model (Figure 1(b)). This scorer comprises a molecule encoder and a spectrum encoder, which evaluate the alignment between molecular structures and spectral data, providing accurate reward scores for solutions during inference. Furthermore, we integrate this scorer into a Monte Carlo Tree Search (MCTS) reasoning framework combined with Self-Refine (Madaan et al., 2023). This framework allows LLMs to optimize previous solutions in a timely manner while balancing the exploration of new solutions and the exploitation of existing ones. By incorporating this scorer into the MCTS framework, LLM can dynamically adjust its reasoning strategy based on feedback, gradually improving the quality of the predicted molecular structure. The scorer also serves as a bridge between the reasoning process and the external knowledge base, using the input spectral data to query and retrieve the most relevant molecular substructures. This effectively mitigates inaccuracies in substructure retrieval, further enhancing the stability and reliability of the model’s reasoning process.

Our contributions can be summarized as follows:

- We propose a **Knowledge-enhanced reasoning framework for Molecular Structure Elucidation: K-MSE**. It applies test-time scaling with Monte Carlo Tree Search and can be integrated with any LLM as a plugin.
- We construct an external molecular substructure knowledge base to supplement LLMs’ chemical knowledge coverage. Additionally, we design a specialized molecule-spectrum scorer as a reward model to accurately evaluate the reasoning outcomes, while also acting as a retriever between the LLM and the knowledge base.
- Experiments on the MolPuzzle benchmark (Guo et al., 2024) demonstrate that our

method significantly enhances performance, specifically achieving an improvement of over 20% on GPT-4o-mini and GPT-4o.

2 Related Works

LLMs for molecular reasoning. LLMs have demonstrated significant potential in the field of chemical molecular research (Zhang et al., 2025; Zhuang et al., 2024). However, directly applying LLMs to molecular-related tasks presents considerable challenges, primarily due to the inherent complexity of molecules and the substantial differences between molecular structures and natural language (Feng et al., 2024; Jiang et al., 2024). Consequently, many studies have sought to enhance the reasoning of LLMs on molecular data by incorporating domain knowledge. For instance, ChemCrow (M. Bran et al., 2024) leverages existing tools to assist LLMs in various downstream tasks; ChatDrug (Liu et al., 2024) integrates domain knowledge into iterative reasoning and feedback loops, enabling LLMs to perform molecular editing tasks; STRUCTCHEM (Ouyang et al., 2024) uses predefined reasoning templates to guide LLMs through complex chemical tasks. However, to the best of our knowledge, no study has yet explored how to leverage LLMs’ reasoning abilities to tackle the challenge of molecular structure elucidation, a crucial problem in chemical experiments that requires deep domain expertise and a precise understanding of molecular structures. The effective application of LLMs to this challenge could significantly enhance the efficiency of scientific experiments and accelerate advancements in chemical research.

LLMs for tree-search reasoning. Multi-path reasoning through tree structures, such as Tree-of-Thought (ToT) (Yao et al., 2023) and Monte Carlo Tree Search (Hao et al., 2023; Chen et al., 2024; Luo et al., 2024; Kang et al., 2024), has proven effective in complex tasks. Some approaches (Qi et al., 2024; Hao et al., 2023) combine problem decomposition with tree search, where each tree expansion step breaks the problem into manageable subproblems. However, these methods necessitate LLM’s intrinsic capability for problem decomposition and a fine-grained reward model to guide the search process. In contrast, other studies (Zhang et al., 2024a,b) minimize reliance on problem decomposition by treating the complete solution as a state within the tree, integrating Self-Refine (Madaan et al., 2023) to gradually expand

the tree structure and optimize reasoning outcomes. Despite their effectiveness, tree search-based reasoning methods highly rely on an accurate reward model to effectively differentiate between desirable and undesirable answers (Uesato et al., 2022). Some researchers (Qi et al., 2024; Zhang et al., 2024a) attempt to leverage the inherent capabilities of LLMs as the reward model, but existing studies have shown that this approach is often limited in its effectiveness (Stechly et al., 2024). We have identified that this issue is particularly pronounced in molecular structure elucidation tasks. In this work, we design a specialized scorer as a reward model to provide more accurate guidance signals, thereby enhancing the reasoning process.

3 Problem Formulation

Molecular structure elucidation refers to determining a molecule’s structure based on various types of spectral data. Following the formulation presented in (Guo et al., 2024), we define the problem as:

$$\hat{m} = f(x_{\text{ir}}, x_{\text{cnmr}}, x_{\text{hnmr}}, x_{\text{formula}}), \quad (1)$$

where \hat{m} represents the predicted molecule in SMILES (Weininger, 1988), x_{ir} is infrared (IR) spectroscopy data, x_{cnmr} and x_{hnmr} are C-NMR and H-NMR spectra, respectively, and x_{formula} denotes molecular formula. The function $f(\cdot)$ represents the reasoning process of LLMs. We denote all input as *question*, where x_{ir} is in image modality, while x_{cnmr} , x_{hnmr} and x_{formula} are in text modality. We provide an illustrative example in Appendix D.

4 Method

In this section, we present the proposed Knowledge-enhanced reasoning framework for Molecular Structure Elucidation (K-MSE). First, we describe the constructed knowledge base, which contains molecular substructures to expand the chemical structure coverage of LLMs. Next, we introduce the molecule-spectrum scorer (Figure 2), a specialized reward model that provides accurate guidance during reasoning and serves as a bridge to retrieve from the knowledge base. Finally, we outline the overall reasoning framework (Figure 3).

4.1 Molecular Substructure Knowledge Base

The molecular structure space is vast and complex, estimated to contain approximately 10^{60} molecules (Polishchuk et al., 2013). This immense space poses a significant challenge for LLMs in

molecular structure prediction, primarily due to the high diversity of molecular structures and the inherent differences between chemical structures and natural language representations (Jiang et al., 2024). LLMs are often adept at recognizing basic molecular structures but struggle to make accurate predictions when faced with more complex situations (exhibited in Figure 1(a)). In other words, LLMs exhibit limitations in achieving a comprehensive coverage of molecular structures.

To address this issue, we construct a molecular substructure knowledge base. Substructures refer to common fragments in molecules that exhibit specific functional or characteristic features, serving as fundamental units for molecular analysis (Fang et al., 2023). This knowledge base contains a broad array of molecular substructures along with their corresponding textual descriptions, which can dynamically provide valuable supplementary information to the LLM during the reasoning process. By leveraging this knowledge base, we equip the LLM with rich, domain-specific chemical knowledge, thereby enhancing its ability to more accurately identify and infer complex molecular structures.

Specifically, the knowledge base \mathcal{KB} can be formulated as:

$$\mathcal{KB} = \{(s_i, d_i)\}, \quad (2)$$

where s_i represents the molecular substructure in SMILES (Weininger, 1988), and d_i is the corresponding textual description. In \mathcal{KB} , both ring and chain substructures are extracted from a widely used molecule database (Polykovskiy et al., 2020), achieving a balance between diversity and universality. By incorporating structural information obtained through external tools as auxiliary input, we use LLM to automate the generation of reliable descriptions for these substructures. The construction details are in Appendix A.1 and Figure 7.

4.2 Molecule-Spectrum Scorer

A reliable reward model is crucial for ensuring that LLMs can effectively correct errors and improve their reasoning capabilities (Uesato et al., 2022). The reward model provides precise feedback signals, guiding LLMs to adjust suboptimal reasoning paths, thus progressively refining the final result. However, as revealed in Figure 1(b), leveraging the inherent capabilities of LLMs as a reward model for molecular structure elucidation tasks has limitations. This is primarily due to the

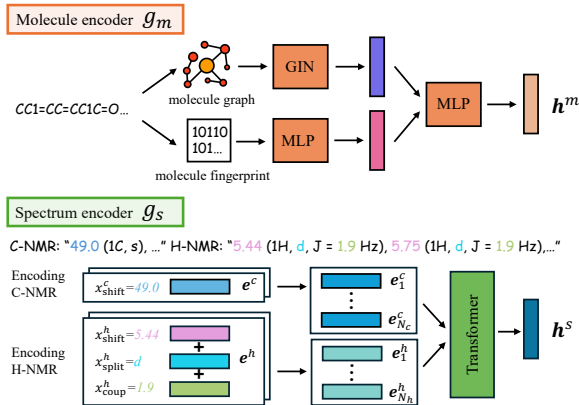


Figure 2: The molecule-spectrum scorer consists of a molecule encoder g_m and a spectrum encoder g_s . After encoding both data types, the reward value is obtained by calculating the similarity between their embeddings h^m and h^s .

lack of deep chemical knowledge in LLMs, which prevents them from accurately assessing the consistency between predicted molecular structures and queried spectral data, making it challenging to differentiate between ideal and non-ideal solutions.

To provide accurate guidance for reasoning, we design a specialized scorer as the reward model. This model consists of a molecule encoder g_m and a spectrum encoder g_s :

$$h^m = g_m(m), \quad h^s = g_s(n), \quad (3)$$

where m is molecule, n is NMR spectrum x_{cnmr} and x_{hnmr} in Equation (1). $h^m, h^s \in \mathbb{R}^d$ are embeddings for molecule and spectrum, respectively.

Typically, the molecule encoder g_m takes both molecule graph m_g (Hu et al., 2020) and fingerprint m_{fp} (Rogers and Hahn, 2010) as input, and employs GIN (Xu et al., 2019) and MLP for encoding:

$$h^m = \text{MLP}(\text{GIN}(m_g), \text{MLP}(m_{fp})). \quad (4)$$

For spectrum, we design a model that leverages the chemical shift of each carbon atom in C-NMR and the chemical shift, splitting pattern, and coupling constants of each hydrogen atom in H-NMR. For example, for a carbon atom with C-NMR spectrum “49.0 (1C, s)”, its chemical shift is 49.0. For a hydrogen atom with H-NMR spectrum “5.44 (1H, d, J = 1.9 Hz)”, its chemical shift, splitting pattern, and coupling constant are 5.44, d , and 1.9, respectively. To represent these features, we discretize all possible values for chemical shifts and coupling constants, assigning a unique token ID to

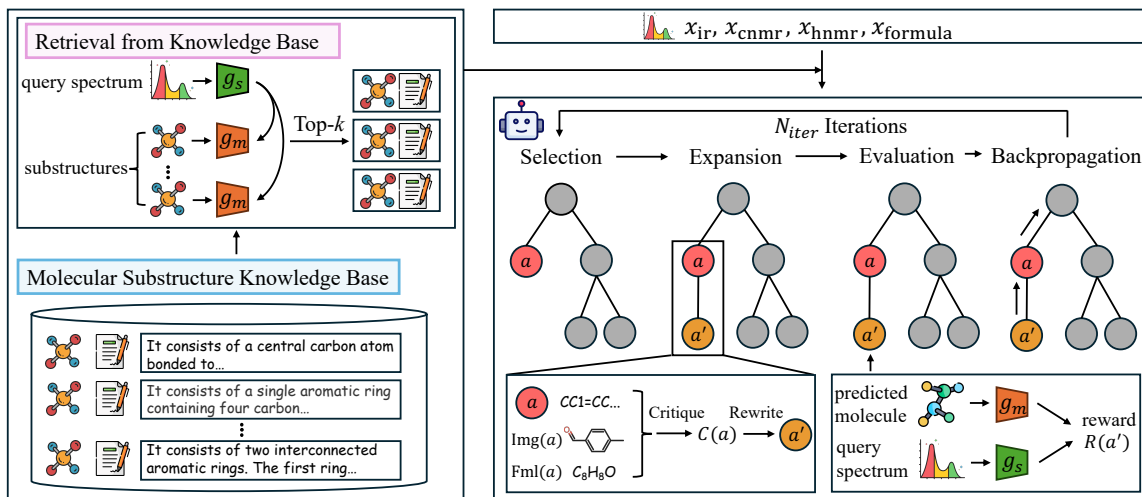


Figure 3: The framework of K-MSE. We construct a molecular substructure knowledge base covering various substructures and their descriptions. For each *question* including $x_{\text{ir}}, x_{\text{cnmr}}, x_{\text{hnmr}}, x_{\text{formula}}$, relevant information is first retrieved from the knowledge base. Then, an MCTS-based iteration is performed, where each node in the tree represents a complete answer. The expansion process critiques previous answers a with image $\text{Img}(a)$ and formula $\text{Fml}(a)$ of the predicted molecule, followed by a rewrite. We derive the reward of answers using the specialized scorer containing g_s and g_m , which is also applied during the retrieval process.

each value. Similarly, we assign a separate token ID to each splitting pattern. The formal representation is as follows:

$$\begin{aligned} e^c &= E(x_{\text{shift}}^c), \\ e^h &= E(x_{\text{shift}}^h) + E(x_{\text{split}}^h) + E(x_{\text{coup}}^h), \end{aligned} \quad (5)$$

where e^c and e^h are embeddings for one carbon and hydrogen atom, respectively. x_{shift}^c is the carbon chemical shift, x_{shift}^h , x_{split}^h and x_{coup}^h are hydrogen chemical shift, splitting pattern and coupling constant. $E(\cdot)$ denotes the embedding function. We then concatenate the embeddings of all carbon and hydrogen atoms and input them into a Transformer encoder (Vaswani et al., 2017):

$$\mathbf{h}^s = \text{Transformer}(e_1^c, \dots, e_{N_c}^c, e_1^h, \dots, e_{N_h}^h), \quad (6)$$

where N_c and N_h are numbers of carbon and hydrogen atoms in this molecule.

For model training, the NT-Xent loss (Chen et al., 2020) is employed:

$$\mathcal{L} = - \sum_{i=1}^B \frac{e^{\text{sim}(\mathbf{h}_i^m, \mathbf{h}_i^s)/\tau}}{\sum_{k=1, k \neq i}^B e^{\text{sim}(\mathbf{h}_i^m, \mathbf{h}_k^s)/\tau} + e^{\text{sim}(\mathbf{h}_i^m, \mathbf{h}_i^s)/\tau}}, \quad (7)$$

where \mathbf{h}_i^m and \mathbf{h}_i^s are molecule and spectrum embeddings for the i -th molecule. $\text{sim}(\cdot, \cdot)$ denotes cosine similarity, B is the batch size, and τ is temperature. Details of the model architecture and training process are in Appendix A.2.

4.3 MCTS-based Reasoning Framework

Effective reasoning requires both knowledge mastery and timely reflection and adjustment. We expand knowledge coverage via a knowledge base and introduce a scorer to evaluate reasoning outcomes accurately. However, LLMs lack self-reflection and dynamic reasoning adjustment, hindering proactive error detection and correction during generation. To address this limitation, we integrate the knowledge base and scorer into a MCTS reasoning framework, which treats each inference process and solution as a tree node (Zhang et al., 2024b), and employs Self-Refine (Madaan et al., 2023) for iterative reflection and improvement during tree search. Prior to MCTS, we first retrieve from the knowledge base. Then MCTS operates iteratively, with each iteration consisting of four key steps: selection, expansion, evaluation, and backpropagation.

Retrieval from Knowledge Base. First, based on the spectrum data in the *question*, we retrieve the most relevant substructures and their corresponding descriptions from the molecular substructure knowledge base. Since the objective of the proposed molecule-spectrum scorer is to measure the correlation between spectral data and molecular structures, it effectively acts as a bridge connecting the query and the knowledge base. Specifically, we use the spectrum encoder g_s to encode the query spectrum and use the molecule encoder

g_m to encode each substructure in the knowledge base. Retrieval is employed as follows:

$$\{(s_j, d_j)\} = \text{Top-}k(\text{sim}(g_m(s_i), g_s(n)),_{s_i \in \mathcal{KB}}), \quad (8)$$

where n is spectrum x_{cnmr} and x_{hnmr} in Equation (1). The retrieved $\{(s_j, d_j)\}$ are then appended after the *question* for subsequent reasoning.

Selection. In the iteration of MCTS, the first step is to select a node for expansion. We use the Upper Confidence Bound applied to Trees (UCT) as selection criterion:

$$a = \arg \max(Q(a_i) + c\sqrt{\frac{\ln N(P(a_i)) + 1}{N(a_i) + \varepsilon}}), \quad (9)$$

where $Q(a_i) \in \mathbb{R}$ is the Q-value of node a_i , which is initialized in the evaluation step and updated in the backpropagation step; $N(\cdot)$ is the number of visits; and $P(\cdot)$ is the parent node. c and ε are hyper-parameters. A node is considered fully selected when its child node number reaches a pre-defined limit.

Expansion. In expansion, a critique process is employed to identify the deficiencies of the current solution in the selected node a :

$$C(a) = \text{Critique}(a, \text{Img}(a), \text{Fml}(a)), \quad (10)$$

where $\text{Img}(a)$ and $\text{Fml}(a)$ represent image and chemical formula of the predicted molecule in a . We observe that for LLMs, the effectiveness of critiquing previous response without providing additional context is significantly limited, as they often struggle to accurately identify molecular structures from text alone. This issue can be mitigated by supplementing with molecular image and chemical formula. Subsequently, a rewrite process is employed to produce a new solution a' as the child of a , based on $C(a)$:

$$a' = \text{Rewrite}(a, C(a)). \quad (11)$$

Evaluation. The newly generated node necessitates an evaluation as its reward value. We employ the proposed molecule-spectrum scorer as the reward model to compute the similarity between the predicted molecular structure and the queried spectrum:

$$R(a') = \text{sim}(g_m(m_{a'}), g_s(n)), \quad (12)$$

where $m_{a'}$ is the predicted molecule in a' , n represents spectrum x_{cnmr} and x_{hnmr} , and $R(a') \in \mathbb{R}$ is the reward value. The Q-value of a' is then initialized with the reward: $Q(a') = R(a')$.

Backpropagation. In the backpropagation step, the Q-value of the new node is propagated to its parent node:

$$Q(a) = 0.5 \times Q(a') + 0.5 \times Q(a). \quad (13)$$

This backpropagation process is iteratively applied until reaching the root. Backpropagation enables the adjustment of node value estimates based on newly acquired feedback, progressively refining the reasoning process and guiding it toward more promising regions of the solution space.

The number of iterations for selection, expansion, evaluation, and backpropagation is controlled by the hyper-parameter N_{iter} . Upon reaching the maximum number of iterations, the node with the highest reward is selected as the final answer.

5 Experiment

5.1 Experimental Setup

We conduct experiments on the MolPuzzle benchmark (Guo et al., 2024), following the ‘‘Addressing Entire Molecule Puzzles’’ setting. IR, C-NMR, H-NMR, and molecular formula serve as inputs, and the goal is to predict the molecule in a *zero-shot* manner. The dataset consists of 216 molecules. We select the following base models: Llama-3.2-11B-Vision-Instruct (Dubey et al., 2024), GPT-4o-mini (gpt-4o-mini-2024-07-18) (OpenAI, 2024a), GPT-4o (gpt-4o-2024-08-06) (OpenAI, 2024b), and GPT-o1 (o1-2024-12-17) (OpenAI, 2024c).

The baseline methods include direct prompting with Chain-of-Thought (CoT) (Wei et al., 2022), as well as employing Self-Refine (Self-R) (Madaan et al., 2023), Self-Consistency (Self-C) (Wang et al., 2023), Multi-Agent Debate (MAD) (Liang et al., 2024), and MCT Self-Refine (MCTSr) (Zhang et al., 2024a). Evaluation metrics include: (1) Fingerprint Tanimoto Similarity (FTS) (Tanimoto, 1958) for chemical similarity between generated and ground-truth molecules using Morgan, MACCS, and RDKit fingerprints; (2) Formula accuracy (Formula Acc), which measures the correctness of chemical formulas of generated molecules; and (3) ACC for the exact structural match. Detailed implementation of the proposed K-MSE and baselines are in Appendix A.3 and B.

5.2 Main Results

Experimental results are presented in Table 1. Our method, K-MSE, significantly outperforms all base-

Model	Morgan FTS	MACCS FTS	RDK FTS	Formula ACC	ACC
Llama-3.2-11B	0.163	0.349	0.188	0.060	0.014
+ Self-R	0.161 (−0.002)	0.345 (−0.004)	0.178 (−0.010)	0.023 (−0.037)	0.023 (+0.009)
+ Self-C	0.172 (+0.009)	0.350 (+0.001)	0.210 (+0.022)	0.032 (−0.028)	0.019 (+0.005)
+ MAD	0.185 (+0.022)	0.368 (+0.019)	0.203 (+0.015)	0.051 (−0.009)	0.009 (−0.005)
+ MCTSr	0.169 (+0.006)	0.343 (−0.006)	0.197 (+0.009)	0.023 (−0.037)	0.004 (−0.010)
+ K-MSE	0.298 (+0.135)	0.465 (+0.116)	0.311 (+0.123)	0.143 (+0.083)	0.111 (+0.097)
GPT-4o-mini	0.260	0.512	0.337	0.185	0.037
+ Self-R	0.287 (+0.027)	0.523 (+0.011)	0.369 (+0.032)	0.231 (+0.046)	0.069 (+0.032)
+ Self-C	0.299 (+0.039)	0.535 (+0.023)	0.373 (+0.036)	0.222 (+0.037)	0.083 (+0.046)
+ MAD	0.292 (+0.032)	0.520 (+0.008)	0.359 (+0.022)	0.222 (+0.037)	0.079 (+0.042)
+ MCTSr	0.281 (+0.021)	0.530 (+0.018)	0.361 (+0.024)	0.176 (−0.009)	0.069 (+0.032)
+ K-MSE	0.470 (+0.210)	0.651 (+0.139)	0.520 (+0.183)	0.412 (+0.227)	0.273 (+0.236)
GPT-4o	0.493	0.690	0.538	0.486	0.278
+ Self-R	0.474 (−0.019)	0.692 (+0.002)	0.522 (−0.016)	0.500 (+0.014)	0.250 (−0.028)
+ Self-C	0.551 (+0.058)	0.732 (+0.042)	0.581 (+0.043)	0.514 (+0.028)	0.347 (+0.069)
+ MAD	0.519 (+0.026)	0.710 (+0.020)	0.558 (+0.020)	0.482 (−0.004)	0.310 (+0.032)
+ MCTSr	0.500 (+0.007)	0.705 (+0.015)	0.551 (+0.013)	0.495 (+0.009)	0.282 (+0.004)
+ K-MSE	0.707 (+0.214)	0.834 (+0.144)	0.727 (+0.189)	0.711 (+0.225)	0.578 (+0.300)
GPT-o1	0.833	0.923	0.856	0.903	0.741
+ K-MSE	0.868 (+0.035)	0.939 (+0.016)	0.888 (+0.032)	0.923 (+0.020)	0.787 (+0.046)

Table 1: Experimental performance on MolPuzzle benchmark. The first row of each group presents the baseline results by directly employing CoT, with the values in () indicating the relative improvement over the baseline. The best results are highlighted in **bold**.

line methods, achieving substantial performance improvements across all base models. Specifically, ACC increases by 0.236 and 0.300 on GPT-4o-mini and GPT-4o, respectively. As the capability of the base models improves, the relative performance gain of K-MSE exhibits an initial increase followed by a decline. We attribute this to that weaker base models (e.g., Llama-3.2-11B) have limited capabilities in chemical reasoning, preventing them from fully exploiting the potential of K-MSE. Conversely, stronger models (e.g., GPT-o1) are closer to the performance ceiling, resulting in diminishing returns with further optimization.

Additionally, some baseline methods lead to performance degradation. For example, MAD and MCTSr cause ACC to decrease by 0.005 and 0.010, respectively, on Llama-3.2-11B. This can be attributed to the insufficient knowledge and reasoning abilities of LLMs in chemical tasks. Both MAD and MCTSr rely on the model’s self-correction mechanism, which requires accurate understanding and adjustment of the reasoning process. Over-reliance on this potentially unreliable mechanism can amplify errors, resulting in greater discrepancies between the generated and ground-truth molecular structures. In contrast, our approach enhances

domain knowledge by incorporating a molecular substructure knowledge base, utilizes specialized scorer to evaluate reasoning outcomes, and introduces additional inputs (such as molecular image and formula) during the self-critique phase. These ensure that K-MSE achieves consistent and significant performance gains across all base models. Further discussion on these can be found in subsequent Sections and Appendix C.

5.3 Detailed Analysis

Specialized Scorer Ensures Precise Answer Evaluation. To evaluate the specialized scorer, we compare it with an LLM-based scorer. The experimental results, shown in Figure 4, highlight the performance differences between the two approaches and analyze the distribution and correlation between the reward score and oracle similarity. The reward score is the score assigned by the scorer to each prediction, while the oracle represents the molecular fingerprint similarity between the predicted and the ground-truth molecules. An ideal scorer should exhibit a positive correlation between these two measures. The results clearly show that the specialized scorer significantly enhances performance, yielding an ACC increase of




Figure 4: Performance comparison (left) and reward vs. oracle similarity (right) between specialized scorer and LLM-based scorer on GPT-4o-mini.




Figure 5: Performance comparison (left) and retrieval hit comparison (right) between scorer as retriever and LLM-based retriever on GPT-4o-mini.

0.139. Notably, the specialized scorer shows a positive correlation of 0.53 between the reward score and the oracle. However, the LLM-based scorer shows a near-zero correlation (0.03) and produces less discriminative, overly lenient scores (generally >0.7), likely due to LLM self-affirmation bias. In contrast, the specialized scorer, as a tailored small model, effectively addresses this bias, providing more accurate and reliable evaluations.

Specialized Scorer Enhances Knowledge Base Retrieval. In K-MSE, the specialized scorer also acts as the retriever between LLM and knowledge base. To assess its effectiveness, we conduct experiments with an LLM-based retriever variant. In this variant, we first prompt the LLM to analyze the molecular structure from the *question* and generate a query description. Then, we use BM25 (Robertson et al., 2009) to compute the similarity between the query description and the substructure descriptions in the knowledge base. Experimental results in Figure 5 demonstrate that the scorer-based retriever significantly outperforms the LLM-based approach, with a 0.120 improvement in ACC and an average 0.169 increase in retrieval hit rate (the proportion of retrieved substructures containing the

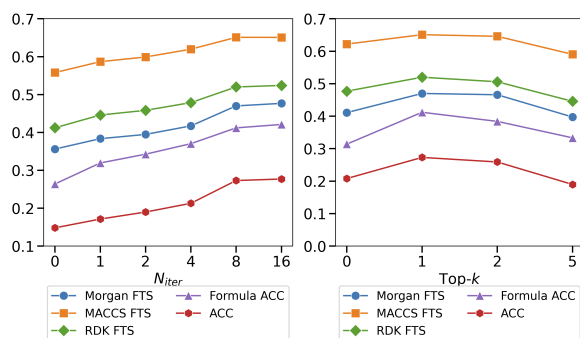


Figure 6: Analysis of MCTS iterations N_{iter} (left) and retrieval counts k from the knowledge base (right).

ground-truth substructures in the target molecule). This highlights the scorer’s ability to retrieve more accurate substructures with fewer retrievals.

MCTS Iterations and Knowledge Base Retrieval Analysis. We analyze the impact of different iteration numbers N_{iter} and retrieval counts k on performance using GPT-4o-mini. Experimental results are shown in Figure 6. Setting N_{iter} being 0 disables the MCTS iterative process, while k being 0 removes the knowledge base. Results show that performance improves with increasing N_{iter} , demonstrating the benefits of iterative reasoning, though gains diminish after a certain point. For retrieval count k , performance initially rises but then declines as k grows, suggesting that while the knowledge base integration is effective, larger k values introduce noise from retrieved irrelevant substructures. Notably, the best performance on GPT-4o-mini occurs at $k = 1$, while on GPT-4o, optimal performance is achieved at $k = 2$. This difference may stem from GPT-4o’s greater robustness, which allows it to more effectively mitigate the impact of retrieval noise.

6 Conclusion

In this work, we introduce K-MSE, a knowledge-enhanced reasoning framework for molecular structure elucidation based on LLMs. K-MSE utilizes a Monte Carlo Tree Search-based test-time scaling approach and can be integrated as a plugin with any LLM. We build a molecular substructure knowledge base and design a specialized molecule-spectrum scorer to address the limited chemical knowledge inherent in LLMs. Experimental results demonstrate the superior performance of K-MSE, and we believe it paves the way for LLMs to function as copilots in the analysis of chemical experiments, thereby accelerating scientific advances.

Limitations

Although the proposed K-MSE framework demonstrates significant performance improvements in molecular structure elucidation, several limitations should be acknowledged. First, the framework operates as a plugin for existing large language models (LLMs) rather than training LLMs specifically for the task. While this design choice enhances flexibility and reduces dependency on task-specific training, the performance of K-MSE is inherently constrained by the capabilities of the underlying LLM. This limitation is particularly pronounced due to the scarcity of reliable and publicly available molecular spectral data, which hinders the development of more specialized models. Second, the evaluation of K-MSE is currently limited to the MolPuzzle benchmark, as it is the only publicly accessible dataset through reliable experimental studies in this domain. While MolPuzzle provides a valuable starting point, its limited scale and diversity may not fully capture the complexity of real-world molecular structure elucidation tasks. To address these limitations, future work will focus on (1) collecting a larger and more diverse dataset to enable more robust training and evaluation, and (2) incorporating cutting-edge LLM training strategies, such as slow-thinking approaches, to further enhance LLM’s reasoning capabilities.

Acknowledgements

This work is funded by the “Pioneer” and “Leading Goose” R&D Program of Zhejiang (Grant No. 2025C01097), NSFCU23B2055, NSFC2302433, NSFCU23A20496, the Fundamental Research Funds for the Central Universities (226-2023-00138), Zhejiang Provincial Natural Science Foundation of China (LQ24F020007) and Hangzhou West Lake Pearl Project Leading Innovative Youth Team Project (TD2023017).

References

Damiano Banfi and Luc Patiny. 2008. www.nmrdb.org: Resurrecting and processing nmr spectra on-line. *Chimia*, 62(4):280–280.

Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. 2024. Step-level value preference optimization for mathematical reasoning. In *EMNLP (Findings)*, pages 7889–7903. Association for Computational Linguistics.

Huajun Chen. 2024. Large knowledge model: Perspectives and challenges. *Data Intelligence*, 6(3):587–620.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.

Tianwei Dai, Sriram Vijayakrishnan, Filip T Szczypinski, Jean-François Ayme, Ehsan Simaei, Thomas Fellowes, Rob Clowes, Lyubomir Kotopantov, Caitlin E Shields, Zhengxue Zhou, et al. 2024. Autonomous mobile robots for exploratory synthetic chemistry. *Nature*, pages 1–8.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Yin Fang, Qiang Zhang, Ningyu Zhang, Zhuo Chen, Xiang Zhuang, Xin Shao, Xiaohui Fan, and Huajun Chen. 2023. Knowledge graph-enhanced molecular contrastive learning with functional prompt. *Nature Machine Intelligence*, 5(5):542–553.

Kehua Feng, Keyan Ding, Weijie Wang, Xiang Zhuang, Zeyuan Wang, Ming Qin, Yu Zhao, Jianhua Yao, Qiang Zhang, and Huajun Chen. 2024. Sciknoweval: Evaluating multi-level scientific knowledge of large language models. *CoRR*, abs/2406.09098.

Leslie D Field, Hsiu Lin Li, and Alison M Magill. 2020. *Organic structures from spectra*. John Wiley & Sons.

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 *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. Reasoning with language model is planning with world model. In *EMNLP*, pages 8154–8173. Association for Computational Linguistics.

Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S. Pande, and Jure Leskovec. 2020. Strategies for pre-training graph neural networks. In *ICLR*. OpenReview.net.

Yinuo Jiang, Xiang Zhuang, Keyan Ding, Qiang Zhang, and Huajun Chen. 2024. Enhancing cross text-molecule learning by self-augmentation. In *ACL (Findings)*, pages 9551–9565. Association for Computational Linguistics.

Jikun Kang, Xin Zhe Li, Xi Chen, Amirreza Kazemi, and Boxing Chen. 2024. Mindstar: Enhancing math reasoning in pre-trained llms at inference time. *CoRR*, abs/2405.16265.

- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *EMNLP*, pages 17889–17904. Association for Computational Linguistics.
- Shengchao Liu, Jiong Xiao Wang, Yijin Yang, Chengpeng Wang, Ling Liu, Hongyu Guo, and Chaowei Xiao. 2024. Conversational drug editing using retrieval and domain feedback. In *ICLR*. OpenReview.net.
- Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, Jiao Sun, and Abhinav Rastogi. 2024. Improve mathematical reasoning in language models by automated process supervision. *CoRR*, abs/2406.06592.
- Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. 2024. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, pages 1–11.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-refine: Iterative refinement with self-feedback. In *NeurIPS*.
- OpenAI. 2023. GPT-4 technical report. *CoRR*, abs/2303.08774.
- OpenAI. 2024a. *Gpt-4o mini: advancing cost-efficient intelligence*. OpenAI.
- OpenAI. 2024b. *Hello gpt-4o*. OpenAI.
- OpenAI. 2024c. *Learning to reason with llms*. OpenAI.
- Siru Ouyang, Zhuosheng Zhang, Bing Yan, Xuan Liu, Yejin Choi, Jiawei Han, and Lianhui Qin. 2024. Structured chemistry reasoning with large language models. In *ICML*. OpenReview.net.
- Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. 2013. Estimation of the size of drug-like chemical space based on gdb-17 data. *Journal of computer-aided molecular design*, 27:675–679.
- Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, et al. 2020. Molecular sets (moses): a benchmarking platform for molecular generation models. *Frontiers in pharmacology*, 11:565644.
- Zhenting Qi, Mingyuan Ma, Jiahang Xu, Li Lyna Zhang, Fan Yang, and Mao Yang. 2024. Mutual reasoning makes smaller llms stronger problem-solvers. *CoRR*, abs/2408.06195.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- David Rogers and Mathew Hahn. 2010. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling LLM test-time compute optimally can be more effective than scaling model parameters. *CoRR*, abs/2408.03314.
- Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. 2024. On the self-verification limitations of large language models on reasoning and planning tasks. *CoRR*, abs/2402.08115.
- Teague Sterling and John J Irwin. 2015. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337.
- Taffee T Tanimoto. 1958. Elementary mathematical theory of classification and prediction.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, H. Francis Song, Noah Y. Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. Solving math word problems with process- and outcome-based feedback. *CoRR*, abs/2211.14275.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Ziyu Wan, Xidong Feng, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun Wang. 2024. Alphazero-like tree-search can guide large language model decoding and training. In *ICML*. OpenReview.net.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *ICLR*. OpenReview.net.
- 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 *NeurIPS*.
- David Weininger. 1988. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks? In *ICLR*. OpenReview.net.

- Xi Xue, Hanyu Sun, Minjian Yang, Xue Liu, Hai-Yu Hu, Yafeng Deng, and Xiaojian Wang. 2023. Advances in the application of artificial intelligence-based spectral data interpretation: A perspective. *Analytical Chemistry*, 95(37):13733–13745.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *NeurIPS*.
- Di Zhang, Xiaoshui Huang, Dongzhan Zhou, Yuqiang Li, and Wanli Ouyang. 2024a. Accessing GPT-4 level mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *CoRR*, abs/2406.07394.
- Di Zhang, Jianbo Wu, Jingdi Lei, Tong Che, Jiatong Li, Tong Xie, Xiaoshui Huang, Shufei Zhang, Marco Pavone, Yuqiang Li, Wanli Ouyang, and Dongzhan Zhou. 2024b. Llama-berry: Pairwise optimization for o1-like olympiad-level mathematical reasoning. *CoRR*, abs/2410.02884.
- Qiang Zhang, Keyan Ding, Tianwen Lv, Xinda Wang, Qingyu Yin, Yiwen Zhang, Jing Yu, Yuhao Wang, Xiaotong Li, Zhuoyi Xiang, et al. 2025. Scientific large language models: A survey on biological & chemical domains. *ACM Computing Surveys*, 57(6):1–38.
- Xiang Zhuang, Keyan Ding, Tianwen Lyu, Yinuo Jiang, Xiaotong Li, Zhuoyi Xiang, Zeyuan Wang, Ming Qin, Kehua Feng, Jike Wang, Qiang Zhang, and Huajun Chen. 2024. Instructbiomol: Advancing biomolecule understanding and design following human instructions. *CoRR*, abs/2410.07919.

A Implementation of K-MSE

A.1 Construction of Molecular Substructure Knowledge Base

The molecular substructures in the knowledge base are divided into two parts: ring structures and chain structures. A ring structure refers to a closed loop of atoms, such as a benzene ring (C1=CC=CC=C1), where the atoms form a cyclic arrangement. A chain structure, on the other hand, refers to a linear arrangement of atoms within the molecule, such as a propyl group (CCC), where the atoms are connected in a straight or branched chain. The extraction of ring structures is defined as follows: first, individual ring structures within the molecule are identified. Then, rings that share two or more atoms are merged. Finally, atoms that are connected to the ring atoms via chemical bonds are included. The extraction of chain structures is defined as follows: first, carbon atoms that are not part of any ring are identified. Then, for each carbon atom, the substructures within its two-hop neighbors are extracted. Finally, substructures containing atoms that are part of a ring are filtered out. We collect substructures from over 4 million molecules in the Moses database (Polykovskiy et al., 2020) using the extraction methods described above. Substructures that appear fewer than 1,000 times are filtered out to maintain a balance between diversity and universality. Ultimately, the knowledge base contains 593 substructures, and substructures are represented in SMILES format.

Another crucial component of the knowledge base is the natural language description of molecular substructures. We use LLMs to automate the process of describing these substructures. Since LLMs may not always accurately recognize molecular structures from SMILES notation, we complement the input by incorporating structural information obtained through external tools. These supplementary structural details include triplet-form structure descriptions, molecular formula, and molecular image. For the triplet-form structure descriptions, we leverage RDKit to extract molecular structure information using the SMILES notation as input. This information includes the number of rings, the size and type of each ring, the types of atoms in the rings, and the functional groups present, such as ether, ester, ketone, imine, and nitrile, etc. Additionally, the corresponding molecular formula and image are also obtained using RDKit. In our implementation, we use GPT-4o-mini to generate natu-

ral language descriptions of molecular structures. Our manual spot checks have shown that LLM can provide accurate structural descriptions by incorporating external tools in this manner. The entire knowledge base construction process is presented in Figure 7, and the prompt used in description generation is in Figure 10.

A.2 Details of Specialized Molecule-Spectrum Scorer

Model architecture. The scorer consists of a molecule encoder g_m and a spectrum encoder g_s , which are responsible for encoding the molecular and spectral data, respectively. The molecule encoder g_m includes a GIN (Graph Isomorphism Network) (Xu et al., 2019) for encoding the molecular graph and an MLP (Multilayer Perceptron) for encoding the molecular fingerprint. The molecular graph can be obtained using RDKit, where the nodes represent atoms and the edges represent chemical bonds. The GIN has a dimension of 300 and 5 layers. The input molecular fingerprint is formed by concatenating Morgan, MACCS, and RDK fingerprints, resulting in a concatenated dimension of 2215, with Morgan and RDK fingerprints having dimensions of 1024 and MACCS having a dimension of 167. The concatenated features are processed by two layers of MLP, with the output dimension being 300. Then, the encoded molecular graph and fingerprint are concatenated, passed through two additional layers of MLP, and ultimately produce a molecular representation h_m with a dimension of 256. In the spectrum encoder g_s , the number of chemical shift tokens for C-NMR is 2,500, while for H-NMR, it is 1,500. Additionally, the number of splitting pattern tokens is 75, and the number of coupling constant tokens is 200. The Transformer encoder (Vaswani et al., 2017) used has a dimension of 256 and consists of 2 layers.

Training. We collect 9,000 molecules from ZINC (Sterling and Irwin, 2015) and use simulation method (Banfi and Patiny, 2008) to obtain the corresponding C-NMR and H-NMR data. Additionally, we gather extra 200 molecules to serve as the validation set. The training is conducted over 100 epochs, with hyperparameter τ set to 0.07. The model checkpoint corresponding to the epoch with the lowest validation loss is saved.

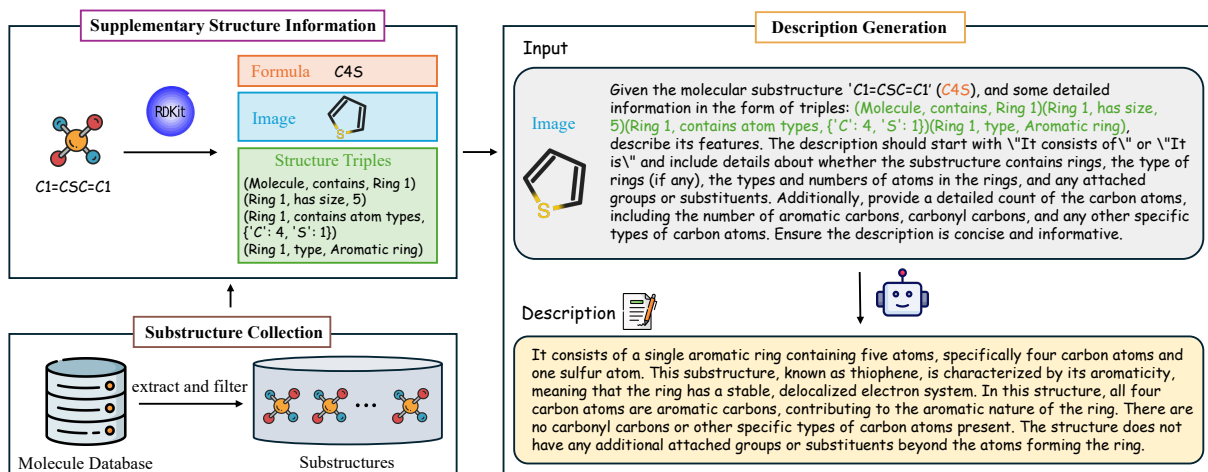


Figure 7: Construction process of the molecular substructure knowledge base. The process consists of three steps: First, collect molecular substructures from a molecule database. Then, RDKit is used to obtain supplementary structure information for each substructure, including molecular formula, structure image, and structure triplets. Finally, LLM is prompted to generate the corresponding textual descriptions.

A.3 Inference Setting

During the inference phase, the parameters c and ϵ in Equation (9) of MCTS are set to 1 and 0.1, respectively, with a maximum of 2 child nodes per node. When the base models are Llama-3.2-11B, GPT-4o-mini, and GPT-4o, N_{iter} is set to 8; when the base model is GPT-o1, N_{iter} is set to 2. For Llama-3.2-11B and GPT-4o-mini, the number of retrieval count k from the knowledge base is set to 1; while for GPT-4o and GPT-o1, k is set to 2. To maximize reproducibility, we set the temperature to 0.0 (most conservative) and top_p to 1.0. The prompts used during inference are shown in Figures 11, 12 and 13.

B Baselines

CoT. Chain-of-Thought (CoT) (Wei et al., 2022) enhances the reasoning capabilities of LLMs by prompting them with “*think step by step*” and generating intermediate reasoning steps to tackle complex tasks.

Self-Refine. Self-Refine (Self-R) (Madaan et al., 2023) iteratively improves the outputs of LLMs by generating feedback and refining its previous outputs without requiring additional training or supervised data. In the experiments, the number of refinement iterations is also set to N_{iter} .

Self-Consistency. Self-Consistency (Self-C) (Wang et al., 2023) improves reasoning in LLMs by sampling multiple reasoning paths and selecting the most consistent answer. In the

experiments, the number of reasoning paths is also set to N_{iter} .

MAD. The Multi-Agent Debate (MAD) (Liang et al., 2024) encourages divergent thinking in LLMs by involving multiple agents in a debate process, where they express their arguments and a judge manages the debate to reach a final solution. In the experiments, the maximum number of debate rounds is 4.

MCTSr. MCT Self-Refine (MCTSr) (Zhang et al., 2024a) integrates LLMs with Monte Carlo Tree Search (MCTS) to enhance performance in complex reasoning tasks by iteratively refining solutions through systematic exploration and heuristic self-evaluation mechanisms. In the experiments, the number of iterations is also set to N_{iter} .

C Additional Experimental Results

C.1 Universal Performance Gains through Knowledge Base Integration

LLMs often face the challenge of insufficient knowledge when handling tasks in the chemical domain. To address this, the proposed molecular substructure knowledge base can be used as a plugin, integrated into any LLM’s reasoning framework. We conduct experiments by integrating this knowledge base with all baseline reasoning methods, and the results are shown in Table 2. The findings indicate that incorporating the knowledge base consistently improves model performance, validating its effectiveness in enhancing reasoning capa-

	Morgan FTS	MACCS FTS	RDKit FTS	Formula ACC	ACC
CoT	0.260	0.512	0.337	0.185	0.037
CoT+ \mathcal{KB}	0.337 (+0.077)	0.572 (+0.060)	0.413 (+0.076)	0.255 (+0.070)	0.111 (+0.074)
Self-R	0.287	0.523	0.369	0.231	0.069
Self-R+ \mathcal{KB}	0.372 (+0.085)	0.598 (+0.075)	0.444 (+0.075)	0.356 (+0.125)	0.157 (+0.088)
Self-C	0.299	0.535	0.373	0.222	0.083
Self-C+ \mathcal{KB}	0.398 (+0.099)	0.602 (+0.067)	0.458 (+0.085)	0.347 (+0.125)	0.171 (+0.088)
MAD	0.292	0.520	0.359	0.222	0.079
MAD+ \mathcal{KB}	0.343 (+0.051)	0.563 (+0.043)	0.403 (+0.044)	0.273 (+0.051)	0.130 (+0.051)
MCTS _r	0.281	0.530	0.361	0.176	0.069
MCTS _r + \mathcal{KB}	0.333 (+0.052)	0.551 (+0.021)	0.414 (+0.053)	0.267 (+0.091)	0.106 (+0.037)

Table 2: Performance of knowledge base integration on GPT-4o-mini across all methods. + \mathcal{KB} denotes the incorporation of the knowledge base. The values in () indicate the relative performance improvement compared to when the knowledge base is not used.

	Morgan FTS	MACCS FTS	RDKit FTS	Formula ACC	ACC
K-MSE	0.470	0.651	0.520	0.412	0.273
w/o Fml	0.432	0.627	0.474	0.347	0.218
w/o Img	0.451	0.628	0.502	0.380	0.236
w/o Fml, w/o Img	0.427	0.620	0.469	0.324	0.208

Table 3: Experimental results of critique input ablation study.

bilities. This also demonstrates that the knowledge base is highly adaptable and flexible, capable of functioning across various reasoning frameworks.

C.2 Ablations on Critique Input

The key to iterative self-reflection and improvement lies in the model’s ability to identify and correct defects in the previous reasoning step. However, we find that due to the limited capability of LLMs in the field of chemical molecules, they often struggle to accurately analyze the molecular structure of the molecule predicted in the previous step. Specifically, when analyzing the predicted molecule represented in SMILES format, the model frequently has difficulty correctly interpreting the corresponding molecular structure during the self-critique process, making the critique unreliable and weakening the effectiveness of self-reflection and improvement.

To address this issue, as shown in Equation (10), we introduce molecular formula and image of the predicted molecule during the critique process, both of which can be generated using RDKit. This additional information helps LLMs more accurately identify the molecular structure, thereby improving the reliability of the critique and enhancing the effectiveness of self-reflection.

We further conduct ablation experiments, and the results are shown in Table 3, where w/o Img and w/o Fml refer to settings in which the molecular image and molecular formula are removed during the critique process, respectively. The results indicate that removing either image or formula leads to a decline in model performance, suggesting that these supplementary inputs play a crucial role in the critique process. They can improve the model’s ability to recognize molecular structures, thereby ensuring the effectiveness of the critique and self-improvement mechanism.

C.3 Ablations on the Component of Specialized Scorer

We analyze the specialized molecule-spectrum scorer through ablation experiments. The variants in the experiments include: w/o m_g (removing molecular graph input in the molecule encoder g_m); w/o m_{fp} (removing molecular fingerprint input in the molecule encoder g_m); w/o x_{split}^h (removing hydrogen atom splitting pattern input in the spectrum encoder g_s); and w/o x_{coup}^h (removing hydrogen atom coupling pattern input in the spectrum encoder g_s). Figure 8 shows the visualization of the validation loss during the training process for these variants, where “all” represents the complete scorer.

Question

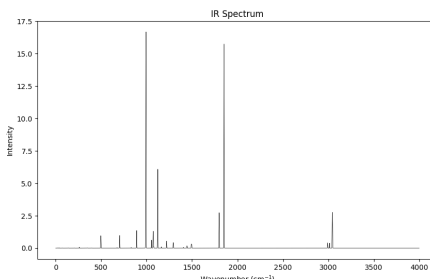
Infrared spectroscopy data (x_{ir})	
C-NMR data (x_{cnmr})	9.1 (2C, s), 27.8 (2C, s), 170.5 (2C, s)
H-NMR data (x_{hnmr})	1.06 (6H, t, J = 7.3 Hz), 2.50 (4H, q, J = 7.3 Hz)
Molecular formula (x_{formula})	C6H10O3
<i>Target</i>	
molecule in SMILES	CCC(=O)OC(=O)CC

Table 4: Example input and target data (ID: 5-43).

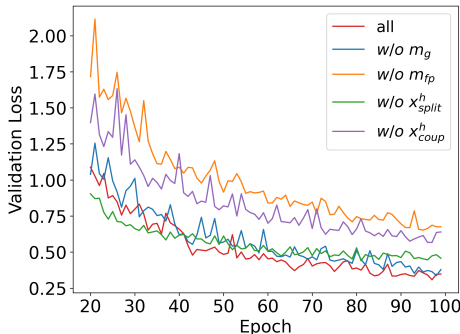


Figure 8: Visualization of validation loss during the scorer training process.

From the results, it is observed that removing any component causes an increase in the validation loss after training convergence. This indicates that each component plays a role in capturing the relationship between the molecule and the spectrum data, and the removal of these components affects the model’s performance.

C.4 Token Consumption Analysis

In Figure 9, we analyze the token consumption on GPT-4o-mini of different methods when solving each problem (left) and the relationship between token consumption and performance (right). K-MSE maintains a relatively reasonable token consumption level, which is lower than that of the Self-R and MCTSr methods. Notably, K-MSE, as an efficient test-time scaling method, effectively utilizes token consumption during test-time to significantly enhance model performance. We attribute this to two key factors: first, the effective expan-




Figure 9: Left: token consumption per *question* for each method. Right: ACC performance versus token consumption across different methods.

sion of the LLM’s knowledge capabilities through the introduction of a molecular substructure knowledge base; second, the development of a specialized scorer for evaluating reasoning outcomes and enhancing self-critique, combined with additional inputs such as molecular images and formulas, enables LLMs to reflect and improve. In contrast, while some methods increase token consumption to acquire more reasoning steps, their trade-off between resource consumption and performance improvement is not optimal.

D Illustrative Data Example

We provide an illustrative data example (ID 5-43 in the dataset) in Table 4 to illustrate the task’s input-output specifications.

Prompt for chain structure

Given the molecular substructure '<##SMILES>' (<##FORMULA>) and some detailed information in the form of triples: <##TRIPLE>, describe its features. The description should start with \"It consists of\" or \"It is\" and include details about the core structure, attached groups or substituents, and any functional groups (e.g., ether, ester, ketone, etc.). Be concise and informative, highlighting key functional groups and their positions where relevant.

Prompt for ring structure

Given the molecular substructure '<##SMILES>' (<##FORMULA>) and some detailed information in the form of triples: <##TRIPLE>, describe its features. The description should start with \"It consists of\" or \"It is\" and include details about whether the substructure contains rings, the type of rings (if any), the types and numbers of atoms in the rings, and any attached groups or substituents. Additionally, provide a detailed count of the carbon atoms, including the number of aromatic carbons, carbonyl carbons, and any other specific types of carbon atoms. Ensure the description is concise and informative.

Figure 10: Prompt for generating descriptions in construction of molecular substructure knowledge base.

You are a chemistry expert and your task is to analyze the data based on IR, C-NMR and H-NMR. Please analyze and think step by step about the structure of the molecule based on the molecular formula and spectrum data. Pay attention to the position of the peaks and the corresponding number of atoms. The final output of the SMILES of the molecule is required to end with 'The SMILES of the molecule is'.

The first input image is IR spectrum.
The molecular formula of the molecule is: '<##FORMULA>'
C-NMR spectra: '<##CNMR-DATA>'
H-NMR spectra: '<##HNMR-DATA>'

You are given a list of molecular substructures and their corresponding natural language descriptions. The molecule could contain one of these substructures, but it may also NOT contain ANY of them. You need to assess each description carefully and decide which substructures are likely to be part of the molecule, considering the overall molecular formula and chemical reasoning.
Input List: <##LIST>

Figure 11: Prompt for initialization.

The provided NMR data is accurate and reliable. You are a strict critic.

Your task is to critically evaluate the predicted molecular formula and structure in current answer and compare them to the target molecular formula and NMR data:

1. Assess whether NMR data analysis has been interpreted correctly, identifying any discrepancies or errors in the analysis.
2. Evaluate whether the predicted molecular formula is consistent with the target molecular formula, noting any deviations or mismatches.
3. Examine the predicted molecular structure in detail to determine whether it fully aligns with the information provided by the NMR data. Highlight any specific substructures that do not match the NMR data analysis criteria. Analyze whether the number of various carbons (aromatic carbons, carbonyl carbons, etc.) or various hydrogens (hydrogens on aromatic carbons, hydrogens on methyl carbons, etc.) in the predicted molecule aligns with the requirements of the NMR data.

You are NOT required to provide corrected prediction. Focus solely on highlighting the areas where the analysis or prediction falls short, and offer a thorough critique of the issues you find.

The first input image is IR spectrum.
The second input image is the molecule predicted in current answer.

```
<problem>
<##PROBLEM>
</problem>

<current_answer>
<##ANSWER>
Predicted Molecule: <##SMILES>
Predicted Molecule Formula: <##pred-FORMULA>
</current_answer>
```

Figure 12: Prompt for critique.

Please refine the current answer based on critique. Ensure that the new answer is different from the current one. Think step by step. The final output of the SMILES of the molecule is required to end with 'The SMILES of the molecule is'.

```
<problem>
<##PROBLEM>
</problem>

<current_answer>
<##ANSWER>
Predicted Molecule: <##SMILES>
Predicted Molecule Formula: <##pred-FORMULA>
</current_answer>

<critique>
<##CRITIQUE>
</critique>
```

Figure 13: Prompt for rewrite.