

MolSight: Optical Chemical Structure Recognition with SMILES Pretraining, Multi-Granularity Learning and Reinforcement Learning

Wenrui Zhang, Xinggang Wang, Bin Feng, Wenyu Liu*

School of Electronic Information and Communications, Huazhong University of Science and Technology

Abstract

Optical Chemical Structure Recognition (OCSR) plays a pivotal role in modern chemical informatics, enabling the automated conversion of chemical structure images from scientific literature, patents, and educational materials into machine-readable molecular representations. This capability is essential for large-scale chemical data mining, drug discovery pipelines, and Large Language Model (LLM) applications in related domains. However, existing OCSR systems face significant challenges in accurately recognizing stereochemical information due to the subtle visual cues that distinguish stereoisomers, such as wedge and dash bonds, ring conformations, and spatial arrangements. To address these challenges, we propose **MolSight**, a comprehensive learning framework for OCSR that employs a three-stage training paradigm. In the first stage, we conduct pre-training on large-scale but noisy datasets to endow the model with fundamental perception capabilities for chemical structure images. In the second stage, we perform multi-granularity fine-tuning using datasets with richer supervisory signals, systematically exploring how auxiliary tasks—specifically chemical bond classification and atom localization—contribute to molecular formula recognition. Finally, we employ reinforcement learning for post-training optimization and introduce a novel stereochemical structure dataset. Remarkably, we find that even with MolSight’s relatively compact parameter size, the Group Relative Policy Optimization (GRPO) algorithm can further enhance the model’s performance on stereomolecular. Through extensive experiments across diverse datasets, our results demonstrate that MolSight achieves state-of-the-art performance in (stereo)chemical optical structure recognition.

Code — <https://github.com/hustvl/MolSight>

Extended version — <https://arxiv.org/abs/2511.17300>

Introduction

A vast amount of chemical information is locked within static images in publications, making it inaccessible for computational analysis. Recent document parsing pipelines (Wang et al. 2024; Li et al. 2025) typically render chemical diagrams as simple images. Optical Chemical Structure

*Corresponding author: liuw@hust.edu.cn
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: **Examples of challenging chemical structures images.** (a) Diversity of images from different sources. (b) 3D molecular information encoded within 2D images.

Recognition (OCSR) addresses this by converting these images into machine-readable formats. This automated process is crucial for accelerating drug discovery and building chemical databases.

General OCSR challenges stem from the diversity of chemical structure sources: (1) Image Degradation: High variance in quality and artifacts from heterogeneous sources like scans and sketches. (2) Symbolic Complexity: The need to parse complex chemical notations beyond the core molecular graph. (3) Stylistic Variance: The lack of standardized drawing conventions, leading to inconsistent representations.

The stereochemical configuration of molecules—including chirality, geometric isomerism, and conformational details—directly impacts their pharmacological effects and chemical behavior, making precise recognition of these structural features a critical requirement for practical applications. Stereochemical recognition introduces additional layers of complexity. As shown in Fig. 1, critical stereochemical indicators such as wedge bonds (representing bonds projecting toward the viewer) and dashed bonds (indicating bonds extending away from the viewer) require precise detection and classification. Furthermore, distinguish-

ing between stereoisomers—molecules with identical connectivity but different spatial arrangements—demands sophisticated understanding of geometric relationships and spatial context that traditional approaches often fail to capture effectively.

To address the aforementioned challenges, we propose MolSight, the first multi-stage learning framework specifically designed for OCSR tasks. In the first stage, MolSight supports pretraining using large quantities of images annotated only with SMILES, with the objective of enhancing the perception capabilities of the image encoder for molecular images. In the second stage, by introducing the additional chemical bond head and coordinate head, the performance of the model can be improved during this process. Recently, reinforcement learning (RL) has shown tremendous potential in improving LLMs’ ability to solve complex reasoning tasks. Inspired by this development, MolSight innovatively introduce RL algorithms into OCSR tasks, utilizing images of stereoisomers that are commonly confused by the model for RL optimization, thereby further enhancing the model’s understanding of chemical semantics.

The main contributions of this paper can be summarized as follows:

- We present MolSight, a comprehensive learning framework for OCSR that enhances model performance across diverse molecular types, particularly stereoisomers, through a three-stage training approach consisting of pre-training, multi-granularity fine-tuning, and RL post-training.
- MolSight represents the first OCSR system to incorporate reinforcement learning methods. By integrating the GRPO algorithm, the model optimization process overcomes the limitations of token-level accuracy and directly optimizes for chemical semantic correctness, effectively improving recognition accuracy for stereoisomeric molecules.
- We construct a new annotated molecular image dataset, Stereo-200k, consisting entirely of challenging stereoisomeric molecules that are prone to confusion. This dataset supports MolSight’s RL training process and will be made publicly available to the research community.
- Extensive experiments demonstrate that MolSight achieves state-of-the-art performance in terms of accuracy, similarity, and robustness, outperforming both classical and learning-based methods across most scenarios, while showing broad potential for downstream applications.

Related Work

Optical Chemical Structure Recognition (OCSR)

Traditional OCSR Methods. Early OCSR systems (McDaniel and Balmuth 1992; Valko and Johnson 2009; Filippov and Nicklaus 2009; Smolov, Zentsev, and Rybalkin 2011; Peryea et al. 2019) were primarily built on handcrafted rules, relying on predefined logic to map image features to chemical structures. These rule-based systems typically employed traditional image processing techniques—including

binarization, denoising, smoothing, and thinning operations—to enhance structural lines and characters. The processed images were then vectorized to convert line elements into vector format, followed by customized optical character recognition (OCR) engines to identify atomic characters. Finally, molecular graph structures were constructed using heuristic methods. Many rule-based OCSR tools continue to receive updates today.

However, these approaches suffer from limited generalization capabilities and struggle to adapt to variations in drawing styles and image quality. Previous studies (Clevert et al. 2021) have demonstrated that even minor perturbations to input images can cause substantial drops in recognition accuracy for these methods.

Deep Learning Approaches. End-to-end deep learning methods (Clevert et al. 2021; Rajan et al. 2023; Qian et al. 2023) typically employ an encoder-decoder architecture that targets line notations of chemical structures (e.g., SMILES (Weininger 1988) and InChI (Heller et al. 2013)) as output. In this framework, image encoders represented by convolutional neural networks (CNNs) are used to extract molecular image features, while decoders based on recurrent neural networks (RNNs) or Transformer generate string-based molecular representations through autoregressive next token prediction. Furthermore, considering the inherent graph properties of molecular structures, previous method (Morin et al. 2023) has integrated graph neural networks (GNNs) into the recognition pipeline. The GNN-based approaches generally require preprocessing steps for component detection or segmentation, and typically face challenges in handling stereochemical molecules and Markush structure molecules.

Reinforcement Learning and Reasoning

The integration of Reinforcement Learning (RL) (Ouyang et al. 2022; Schulman et al. 2017) and Chain-of-Thought (CoT) (Wei et al. 2022) have proven effective in enhancing Large Language Models (LLMs) reasoning performance. CoT prompting enables LLMs to decompose complex problems into intermediate reasoning steps, making the problem-solving process more transparent and systematic. While RL can help LLMs to recognize which reasoning processes are correct through reward-based optimization. Popular methods include Direct Preference Optimization (DPO) (Rafailov et al. 2023), which has been widely adopted for training reasoning-focused models. More recently, advanced techniques such as GRPO (Shao et al. 2024) have been introduced, which lead to more stable and effective training dynamics.

Method

OCSR-Specific Image Captioning Pre-training

Most existing public OCSR datasets contain only chemical structure images and their corresponding SMILES annotations. To leverage these data, a natural approach is to treat the SMILES text as a special description of the chemical structure image, thereby utilizing image captioning techniques to guide model learning. We selected MolParser-7M

(Fang et al. 2024) as our pre-training dataset due to its substantial scale and comprehensive coverage of diverse data sources and image styles. However, the inherent limitations of the original SMILES notation prevent it from describing the numerous Markush structures present in this dataset. To address this limitation, we implemented a simple extension to the SMILES notation system, named SMILES-M. Fig. 2 illustrates the different types of structural variations in Markush structures and their corresponding extended representation methods.

Towards More Powerful OCSR Model through Multi-Granularity Learning

MolScribe (Qian et al. 2023) introduces two additional tasks alongside autoregressive SMILES text generation: chemical bond classification and atom localization. In MolScribe, these three tasks are jointly optimized in a collaborative way. Its experimental results demonstrate that by using the predicted types of chemical bonds and the predicted coordinates of atoms to refine the final molecular structures, the accuracy of OCSR can be improved. We further investigated the impact of these two auxiliary tasks on the OCSR task.

Atoms as Queries After processing SMILES text through the tokenizer, the resulting tokens can be categorized into atom tokens and non-atom tokens. Atom tokens correspond to atoms or functional groups within the molecule, while non-atom tokens represent auxiliary information such as brackets indicating the start and end positions of groups, numbers denoting ring connections, and symbols representing chemical bonds or charges. The queries corresponding to atom tokens contain all the information required for the model to predict the respective atoms. We use these queries as inputs for both the chemical bond head and the coordinate head, as shown in Fig. 3.

Chemical Bond Classification To predict the chemical bond b from the i -th atom a_i to the j -th atom a_j , following MolScribe, we simply concatenate the hidden states of these two atom queries and feed them into a chemical bond head for classification, as:

$$\hat{p}_{i \rightarrow j}(\mathbf{b}) = \phi_{bond}(\mathbf{h}_i \oplus \mathbf{h}_j),$$

where \mathbf{h}_i and \mathbf{h}_j are the hidden states of a_i and a_j , respectively. \oplus indicates the operation of vector vertical concatenation. The chemical bond head ϕ_{bond} is implemented as a 2-layer MLP network followed by a softmax.

In practice, we find that SMILES text generation and chemical bond classification exhibit strong synergy, and jointly optimizing these two tasks can effectively improve the training performance of both. More detailed quantitative comparisons can be found in the Main Results and Ablation Study section.

Atom Localization Following SimCC (Li et al. 2022), we treat the coordinate prediction problem as a one-dimensional classification task for keypoints in both horizontal and vertical directions, which alleviates the localization accuracy issues caused by low-resolution feature maps while maintaining processing speed. MolScribe explicitly inserts coordinate tokens into SMILES sequences for autoregressive

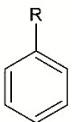
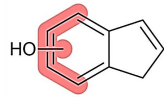
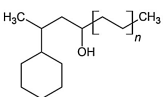
Type	(a) Type changes	(b) Location changes	(c) Frequency changes
Image			
SMILES	⊗ unable to express		
SMILES-M	<chem>[R]c1ccccc1</chem>	<chem>C1=Cc2ccccc2C1<sep>2:OH</chem>	<chem>C[CC?n]C(O)CC(C)C1CCCC1</chem>

Figure 2: **Examples of how SMILES-M express Markush structures.** Our SMILES-M can deal with all types of Markush structures, including type changes, location changes, and frequency changes.

learning, but this approach has two significant limitations: (1) As a classification problem, the correct coordinate label is one-hot encoded, and cross-entropy loss is utilized, meaning all incorrect coordinates receive equal penalties except for the single correct coordinate point. However, it would be more reasonable for model predictions closer to the correct coordinates to receive lower penalties. (2) This method increases the target SMILES sequence length by nearly three-fold, making it more challenging for the model to handle long sequences. Our MolSight continues to extract coordinate information based on atomic queries, achieving higher localization accuracy without increasing sequence length.

To address the first problem mentioned above, it is necessary to introduce a loss function suitable for coordinate classification tasks. Following RTMO (Lu et al. 2024), we employ maximum likelihood estimation (MLE) to optimize the parameters θ of the coordinate head. The true coordinate μ is assumed to follow a Laplace distribution centered at the annotated values μ_g . We use the negative log-likelihood function as our loss function, for the k -th atom:

$$\begin{aligned} \mathcal{L}_{coord} &= -\log [P(\mu = \mu_g | \theta)] \\ &= -\log \left[\sum_{i=1}^B P(\mu_g | x_i, \theta) \hat{p}_k(x_i | \theta) \right], \end{aligned}$$

where

$$\begin{aligned} P(\mu_g | x_i, \theta) &= \frac{1}{2\hat{b}} e^{-\frac{|\mu_g - x_i|}{\hat{b}}}, \\ \hat{p}_k(x_i | \theta) &= \frac{e^{\mathbf{h}_k \cdot \phi_{coord}(\mathbf{PE}(x_i))}}{\sum_{j=1}^B e^{\mathbf{h}_k \cdot \phi_{coord}(\mathbf{PE}(x_j))}}, \end{aligned}$$

the scale parameter \hat{b} of the Laplace distribution is additionally predicted by the coordinate head, it reflects the model’s uncertainty regarding the current coordinate prediction results, and can serve as a reference for filtering subsequent molecular structure prediction results. The positional encoding function \mathbf{PE} is the same as the one from Transformer (Vaswani et al. 2017), B is a hyperparameter denotes the number of coordinate bins, and x_i and x_j are discrete horizontal/vertical coordinate values picked from $linspace(0, 1, B)$.

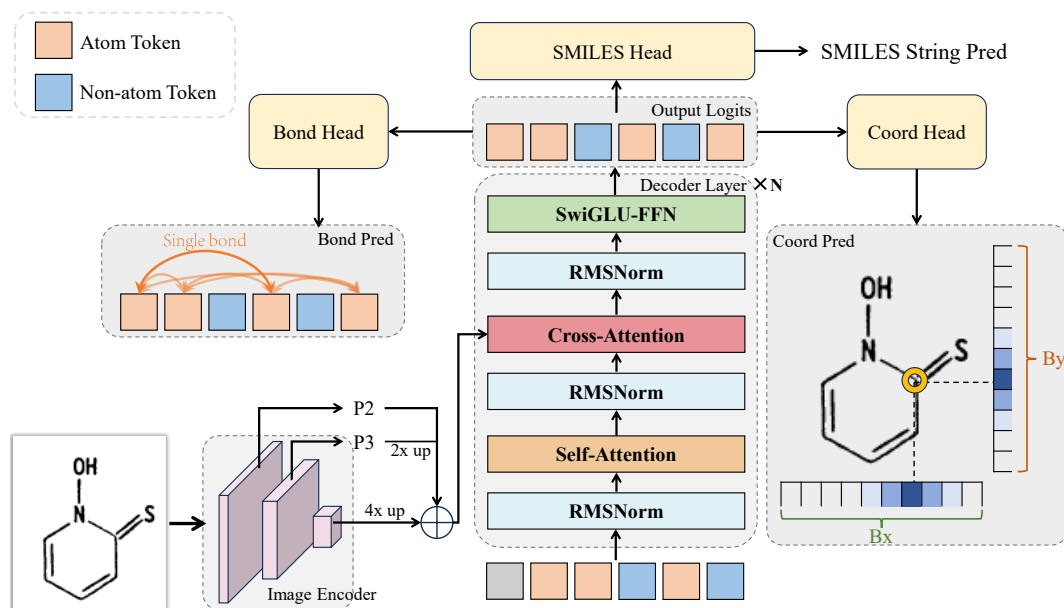


Figure 3: **Overall pipeline of Molsight.** Given a chemical structure image, the image encoder extracts and fuses multi-level image features, which will be fed into the SMILES decoder with previous SMILES tokens to predict the next SMILES token. The SMILES head maps the output logits into SMILES vocabulary space, while the bond head and the coord head predict the chemical bond type and location of each atom token, respectively. The residual connections are omitted in this figure.

In our early experiments, we found that incorporating the atom localization task into our training pipeline negatively impacted the optimization of the other two tasks. Additionally, directly using the hidden states from the SMILES decoder for coordinate prediction failed to achieve satisfactory accuracy. To address these issues, we introduced two additional Transformer layers to process the atom queries before feeding them to the coordinate head. Furthermore, the coordinate head is optimized independently from the other two tasks.

Semantic-aware Reinforcement Learning

SMILES text follows explicit syntactic rules (such as parentheses and ring number pairing), and generation errors are costly—a single incorrect character can cause the entire SMILES invalid, leading to failure in downstream tasks (e.g., preventing use in chemical modeling). Therefore, the metric of primary concern is typically not token-level accuracy, but rather the overall accuracy of SMILES strings (i.e., whether they are completely correct), which is not differentiable and cannot be directly used for gradient optimization. Furthermore, a molecule may have multiple valid SMILES texts, particularly for stereoisomers, where aromatic rings, cis-trans configurations, and chiral centers can be expressed in various correct formats. In such cases, relying solely on imitation learning has inherent limitations, as inconsistent annotation information leads to conflicting optimization objectives.

As shown in Fig. 4, a more direct approach is to allow the model to explore more possible SMILES texts and then apply RL to directly optimize the overall correctness/similarity

Algorithm 1: Reward Calculation for Molecular Structure

Input: Pred SMILES S_{pred} , Ground Truth SMILES S_{gt}
Parameter: Tanimoto weight $w_t = 0.4$, Stereochemistry weight $w_s = 0.6$

Output: Total Reward R_{total}

```

1:  $S'_{pred} \leftarrow \text{ExpandAbbreviations}(S_{pred})$ 
2:  $S'_{gt} \leftarrow \text{ExpandAbbreviations}(S_{gt})$ 
3: Initialize  $R_{total} \leftarrow \emptyset$ 
4: for each pair  $(s_p, s_g) \in (S'_{pred}, S'_{gt})$  do
5:    $r_{tanimoto} \leftarrow 0.0$ 
6:    $r_{stereo} \leftarrow 0.0$ 
7:    $M_p \leftarrow \text{MoleculeFromSMILES}(s_p)$ 
8:    $M_g \leftarrow \text{MoleculeFromSMILES}(s_g)$ 
9:   if  $M_p$  is valid AND  $M_g$  is valid then
10:     $\triangleright$  - Tanimoto Similarity Reward -
11:     $FP_p \leftarrow \text{Fingerprint}(M_p)$ 
12:     $FP_g \leftarrow \text{Fingerprint}(M_g)$ 
13:     $r_{tanimoto} \leftarrow \text{TanimotoSimilarity}(FP_p, FP_g)$ 
14:     $\triangleright$  - Stereochemistry Reward -
15:    if  $\text{InChIKey}(M_p) = \text{InChIKey}(M_g)$  then
16:       $r_{stereo} \leftarrow 1.0$   $\triangleright$  Exact match
17:    else if  $\text{AtomCount}(M_p) = \text{AtomCount}(M_g)$  then
18:       $r_{stereo} \leftarrow 0.3$   $\triangleright$  Similar structure
19:    else
20:       $r_{stereo} \leftarrow 0.1$ 
21:    end if
22:  end if
23:   $r_{combined} \leftarrow w_t \times r_{tanimoto} + w_s \times r_{stereo}$ 
24:  Append  $r_{combined}$  to  $R_{total}$ 
25: end for
26: return  $R_{total}$ 

```

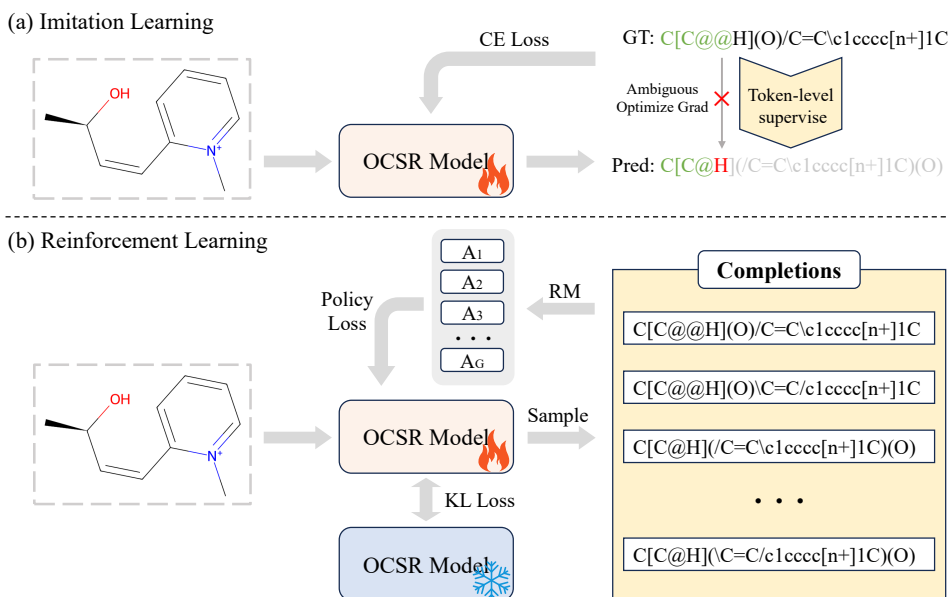


Figure 4: **Comparison of Training Paradigms.** (a) **Imitation Learning:** due to the diversity of correct SMILES text, token-level optimization may cause ambiguous optimize direction. (b) **Reinforcement Learning:** multiple completions are sampled at once, then scored on Tanimoto similarity and structural consistency. We use GRPO to achieve this trajectory-level optimization.

metrics. For stereoisomers, this direct optimization of chemical semantic correctness can also guide the model to attend to stereochemical markers, bond orientations, and other critical details in molecular images, thereby improving the model’s ability to recognize stereoisomeric molecules.

We choose GRPO as the RL algorithm for MolSight, the loss function is defined as follows:

$$\mathcal{L}_{\text{GRPO}}(\theta) = -\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left[\frac{\pi_{\theta}(o_{i,t} | o_{i,<t})}{[\pi_{\theta}(o_{i,t} | o_{i,<t})]_{\text{no grad}}} \hat{A}_{i,t} - \beta \mathbb{D}_{\text{KL}}[\pi_{\theta} \| \pi_{\text{ref}}] \right],$$

where o_1, o_2, \dots, o_G are a group of sampled outputs, the advantage $\hat{A}_{i,t}$ is computed using the normalized reward within the group.

The reward of SMILES consists of two parts: tanimoto similarity reward and stereochem reward. The processing of chemical information is done by RDKit (Landrum, G 2016). The details of our reward function are shown in Alg. 1.

Experiments

Implementation Details

MolSight employs an encoder-decoder architecture. We select EfficientViT (L1 version with approximately 53M parameters) (Cai et al. 2023) as the image encoder due to its lightweight convolutional architecture and linear attention mechanism. The decoder adopts a 6-layer standard Transformer architecture with several common improvements including RoPE (Su et al. 2024), SwiGLU (Shazeer 2020), and RMSNorm (Zhang and Sennrich 2019). The decoder parameters are randomly initialized rather than using pre-trained language models for initialization, as there exists a signifi-

cant vocabulary mismatch between standard language models and SMILES text.

First, we pre-train the model for 2 epochs on a dataset containing only SMILES annotations using next token prediction. Subsequently, we jointly train the model for 10 epochs on both SMILES text generation and chemical bond classification (experimental results demonstrate that extending training to 30 epochs can further improve model performance), while adapting the model with a coordinate head. Finally, we further optimize the model for 2 epochs using RL methods.

Dataset and Evaluation metric

Dataset During the pre-training stage, we train our model on the MolParser-7M dataset from (Fang et al. 2024) for SMILES text generation tasks. In the fine-tuning phase, we employ the PubChem-1M and USPTO-680K from (Qian et al. 2023) for multi-granularity learning. During the post-training phase, we enhance the model using reinforcement learning on our self-collected stereochemical molecular dataset, named Stereo-200K. For evaluation, we report model performance on four classical real-world datasets (USPTO (Filippov and Nicklaus 2009), Maybridge UoB (Sadawi, Sexton, and Sorge 2012), CLEF-2012 (Piroi et al. 2010), and JPO (Fujiyoshi, Nakagawa, and Suzuki 2011)) and four synthetic datasets (Staker (Staker et al. 2019), ChemDraw, Indigo, and Stereo-2K). Both the Stereo-200K and Stereo-2K datasets are introduced for the first time in this work.

Metric For molecular recognition performance, we use exact match accuracy and average molecular fingerprint similarity (i.e., Tanimoto Coefficient) as evaluation metrics. For

Method	USPTO				UoB			CLEF				JPO		
	graph	stereo	exact	tani	graph	exact	tani	graph	stereo	exact	tani	graph	exact	tani
OSRA-2.2.1	86.5	79.0	83.5	92.9	76.1	75.5	88.4	86.2	89.3	83.3	88.4	51.8	50.7	69.1
MolVec-0.9.8	91.4	<u>82.8</u>	88.1	95.8	80.7	80.2	91.5	84.1	<u>89.1</u>	82.8	88.5	68.9	<u>66.4</u>	85.0
Imago-2.0	89.2	77.4	87.2	94.5	58.1	57.6	77.1	65.1	45.1	60.3	84.1	40.7	40.4	59.6
DECIMER-2.7.1	61.5	43.8	58.4	92.4	87.1	86.4	96.2	80.4	76.0	73.8	89.1	40.4	39.3	86.4
MolGrapher†	82.9	10.3	65.7	93.3	84.9	84.2	93.6	71.1	10.4	52.5	86.7	61.1	52.2	76.3
MolScribe	94.6	69.0	<u>88.4</u>	97.5	88.2	87.4	<u>96.0</u>	<u>91.8</u>	76.8	85.5	90.5	60.0	57.6	<u>88.3</u>
MolSight (ours)	<u>94.0</u>	85.1	92.0	<u>97.4</u>	<u>87.9</u>	<u>87.1</u>	95.9	92.2	76.0	<u>84.9</u>	<u>90.1</u>	<u>68.7</u>	66.7	90.7

Table 1: **Performance comparison with existing SOTA methods on real data benchmarks.** †: the stereo-capable variant of MolGrapher was used, **graph**: recognition accuracy ignoring stereoisomers, **stereo**: recognition accuracy on stereochemical molecules, **exact**: exact recognition accuracy, **tani**: Tanimoto Coefficient.

Setting	graph	stereo
<i>Get stereo from SMILES</i>		
SMILES	93.7	80.6
+ edge	94.2	80.9
<i>Get stereo from edges & coords</i>		
Joint Training	43.8	10.6
Separate Training	94.2	29.4
+ extra decoder layers		
w/ L1 Loss	94.2	82.4
w/ MLE Loss	94.2	83.7

Table 2: **Ablation on proposed training tasks.** The results were obtained from USPTO dataset.

atom localization, we employ the Object Keypoint Similarity (OKS) metric to assess the performance, as:

$$OKS = \frac{1}{N_{kpts}} \sum_{n=1}^{N_{kpts}} e^{-\frac{d_n^2}{2s^2}},$$

where N_{kpts} is the number of keypoints, d_n is the Euclidean distance between the n -th predicted location and the ground-truth location, and s is a scale factor.

Main Results and Ablation Study

Experiments on Real Data Benchmarks Our first objective is to develop a plain OCSR model that achieves advanced performance on in-domain data by generating SMILES text alone, thereby facilitating efficient fine-tuning for users on specific target datasets.

Tab. 1 presents a quantitative comparison of our method against current SOTA approaches on four classic real-world benchmarks. The stereochemical recognition accuracy for UoB and JPO datasets is not available since these datasets contain virtually no stereochemical molecules. For fair comparison, both our method and MolScribe were fine-tuned for

Training		Stereo-2k			CLEF		
Strategy		graph	stereo	tani	graph	stereo	tani
SFT		95.4	80.1	98.5	92.2	71.0	90.0
SFT +RL	Tanimoto	97.2	81.8	99.2	92.2	74.3	90.1
	Stereo	96.6	86.5	98.9	92.6	78.4	90.0
	Weighted	96.9	87.1	99.3	91.8	80.6	90.0

Table 3: **Ablation on different training strategies.** Further fine-tuning the model with RL led to additional performance gains. **Tanimoto**: only use Tanimoto Similarity Reward, **Stereo**: only use Stereochemistry Reward, **Weighted**: weighted reward function as shown in Alg. 1.

30 epochs on the same datasets without using chemical bond information or atomic coordinate information for auxiliary correction.

The results demonstrate that compared to MolScribe, our MolSight model exhibits superior performance on in-domain data (i.e., the USPTO dataset), particularly achieving a stereochemical recognition accuracy of 85.1% (+16.1%). Additionally, on the more challenging JPO dataset, MolSight attains an exact recognition accuracy of 66.7% (+9.1%) and a graph recognition accuracy of 68.7% (+8.7%). Furthermore, Our MolSight is the only method that achieves Tanimoto Coefficients exceeding 90% across all four benchmarks.

Ablation Study on Training Tasks In the second training stage, we introduce two auxiliary learning objectives: chemical bond classification and atom localization. This serves a dual purpose. First, we aim to validate whether adding more supervision information can improve the results of SMILES text generation. Second, these auxiliary tasks provide a novel basis for determining the molecule’s stereochemistry. By leveraging the RDKit toolkit, we can infer the complete 3D molecular conformation, including specific

Methods	BBBP	Tox21	ToxCast	SIDER	ClinTox	MUV	HIV	Bace
<i>Multimodal</i>								
3D InfoMax (Stärk et al. 2022)	69.1±1.0	74.5±0.7	64.4±0.8	60.6±0.7	79.9±3.4	74.4±2.4	76.1±1.3	79.7±1.5
GraphMVP (Liu et al. 2021)	68.5±0.2	74.5±0.4	62.7±0.1	62.3±1.6	79.0±2.5	75.0±1.4	74.8±1.4	76.8±1.1
MoleBlend (Yu et al. 2024)	73.0±0.8	77.8±0.8	66.1±0.0	64.9±0.2	87.6±0.7	77.2±2.3	79.0±0.8	83.7±1.4
<i>2D graph</i>								
MolCLR (Wang et al. 2022)	66.6±1.8	73.0±0.1	62.9±0.3	57.5±1.7	86.1±0.9	72.5±2.3	76.2±1.5	71.5±3.1
GraphMAE (Hou et al. 2022)	72.0±0.6	75.5±0.6	64.1±0.3	60.3±1.1	82.3±1.2	76.3±2.4	77.2±1.0	83.1±0.9
Mole-BERT (Xia et al. 2023)	71.9±1.6	76.8±0.5	64.3±0.2	62.8±1.1	78.9±3.0	78.6±1.8	78.2±0.8	80.8±1.4
<i>Image</i>								
EfficientViT (ImageNet pretrain)	63.6±0.8	71.3±0.5	63.2±0.4	60.4±0.4	98.0±0.3	66.9±1.2	73.8±1.2	73.9±1.5
EfficientViT (MolSight pretrain)	68.0±0.6	75.0±0.4	65.0±0.2	62.6±0.7	98.4±0.9	73.6±1.0	74.1±1.1	76.9±1.6

Table 4: **Results on molecular property prediction tasks.** ROC-AUC scores are reported (higher is better). The MolSight pre-training method demonstrates superior performance over the ImageNet (Deng et al. 2009) pre-training approach across all evaluated datasets, and comparable to those advanced molecular pre-training methods that are based on 2D graph or multimodal.

chirality and cis-trans isomerism, from the predicted 2D coordinates and chemical bond types.

Tab. 2 presents the ablation study of our proposed training tasks. The results indicate that augmenting the SMILES decoder with a chemical bond head not only enables the model to predict chemical bonds but also slightly improves the accuracy of SMILES text generation.

However, the introduction of the atom localization task complicates the overall optimization. Jointly optimizing all three objectives causes the model to over-prioritize the atom position detection branch, which in turn hinders the convergence of the primary SMILES generation task. This is reflected in the training metrics as a high OKS score but a low token-level SMILES accuracy.

To address this issue, we decoupled the coordinate prediction branch and optimized it separately. This approach effectively prevents the coordinate prediction from interfering with the optimization of other tasks. Nevertheless, directly using the output queries from the SMILES decoder for coordinate prediction did not yield satisfactory localization performance, so we added two additional decoder layers into the coordinate prediction branch to further extract positional information from the input image. This modification led to a significant improvement, where the accuracy of the molecular stereochemistry inferred from the predicted chemical bonds and coordinates surpassed that derived directly from the generated SMILES text.

Impact of Reinforcement Learning We further investigated the use of RL post-training to encourage the model to autonomously explore correct SMILES notations, thereby enhancing its accuracy in recognizing stereoisomers. To this end, we constructed the Stereo dataset, which is composed entirely of stereoisomer images. These molecular data were sourced from the PubChem database and include a balanced distribution of chiral and cis-trans isomers, as well as a significant number of structurally similar molecules. We utilized the RDKit toolkit to generate the molecular images with random stylistic variations. Tab. 3 presents a compar-

ison of the model’s performance on both in-domain (Stereo-2k) and out-of-domain (CLEF-2012) data, before and after RL-based training on the Stereo-200k dataset.

Analysis of Transfer Learning Performance To validate the transferability of features learned by MolSight, we evaluated our model on downstream tasks from the MoleculeNet (Wu et al. 2018), one of the most widely used benchmarks for molecular property prediction. Our evaluation procedure was plain: First, each molecule was rendered into a 2D image using the Indigo toolkit (Pavlov et al. 2011). Next, the MolSight image encoder, with its parameters frozen, was employed to extract features from each image. These image features, after undergoing global average pooling, were then fed into a trainable MLP probe for the final classification. Following standard evaluation protocols for this benchmark, all datasets were partitioned using the scaffold split method, which groups molecules by their structural backbone to provide a more rigorous test of model generalization. Tab. 4 reports the ROC-AUC scores on all eight classification tasks. We report the mean and standard deviation across 5 random seeds (from 0 to 4).

Conclusion

In this paper, we introduced MolSight, a comprehensive framework designed to address the challenges of Optical Chemical Structure Recognition (OCSR), particularly in accurately interpreting stereochemistry. Our novel three-stage training paradigm, which synergizes large-scale pre-training, multi-granularity fine-tuning, and reinforcement learning, proves highly effective. Extensive experiments demonstrate that MolSight achieves state-of-the-art performance across several challenging benchmarks. Furthermore, the strong transfer learning performance on downstream molecular property prediction tasks validates the robustness of the features learned by our model, highlighting its potential for broader applications in automated chemical data analysis.

Acknowledgments

This work was partially supported by the National Science and Technology Major Project under Grant No. 2023YFF0905400.

References

- Cai, H.; Li, J.; Hu, M.; Gan, C.; and Han, S. 2023. Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17302–17313.
- Clevert, D.-A.; Le, T.; Winter, R.; and Montanari, F. 2021. Img2Mol—accurate SMILES recognition from molecular graphical depictions. *Chemical science*, 12(42): 14174–14181.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Fang, X.; Wang, J.; Cai, X.; Chen, S.; Yang, S.; Tao, H.; Wang, N.; Yao, L.; Zhang, L.; and Ke, G. 2024. Molparser: End-to-end visual recognition of molecule structures in the wild. *arXiv preprint arXiv:2411.11098*.
- Filippov, I. V.; and Nicklaus, M. C. 2009. Optical structure recognition software to recover chemical information: OSRA, an open source solution.
- Fujiyoshi, A.; Nakagawa, K.; and Suzuki, M. 2011. Robust method of segmentation and recognition of chemical structure images in cheminfy. In *Pre-proceedings of the 9th IAPR international workshop on graphics recognition, GREC*, volume 1.
- Heller, S.; McNaught, A.; Stein, S.; Tchekhovskoi, D.; and Pletnev, I. 2013. InChI-the worldwide chemical structure identifier standard. *Journal of cheminformatics*, 5(1): 7.
- Hou, Z.; Liu, X.; Cen, Y.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. GraphMAE: Self-Supervised Masked Graph Autoencoders. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 594–604.
- Landrum, G. 2016. Rdkit: open-source cheminformatics. <http://www.rdkit.org>.
- Li, Y.; Yang, S.; Liu, P.; Zhang, S.; Wang, Y.; Wang, Z.; Yang, W.; and Xia, S.-T. 2022. Simcc: A simple coordinate classification perspective for human pose estimation. In *European Conference on Computer Vision*, 89–106. Springer.
- Li, Z.; Liu, Y.; Liu, Q.; Ma, Z.; Zhang, Z.; Zhang, S.; Guo, Z.; Zhang, J.; Wang, X.; and Bai, X. 2025. MonkeyOCR: Document Parsing with a Structure-Recognition-Relation Triplet Paradigm. *arXiv:2506.05218*.
- Liu, S.; Wang, H.; Liu, W.; Lasenby, J.; Guo, H.; and Tang, J. 2021. Pre-training molecular graph representation with 3d geometry. *arXiv preprint arXiv:2110.07728*.
- Lu, P.; Jiang, T.; Li, Y.; Li, X.; Chen, K.; and Yang, W. 2024. Rtm0: Towards high-performance one-stage real-time multi-person pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1491–1500.
- McDaniel, J. R.; and Balmuth, J. R. 1992. Kekule: OCR-optical chemical (structure) recognition. *Journal of chemical information and computer sciences*, 32(4): 373–378.
- Morin, L.; Danelljan, M.; Agea, M. I.; Nassar, A.; Weber, V.; Meijer, I.; Staar, P.; and Yu, F. 2023. MolGrapher: Graph-based Visual Recognition of Chemical Structures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 19552–19561.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pavlov, D.; Rybalkin, M.; Karulin, B.; Kozhevnikov, M.; Savelyev, A.; and Churinov, A. 2011. Indigo: universal cheminformatics API. *Journal of cheminformatics*, 3(Suppl 1): P4.
- Peryea, T.; Katzel, D.; Zhao, T.; Southall, N.; and Nguyen, D.-T. 2019. MOLVEC: Open source library for chemical structure recognition. In *Abstracts of papers of the American Chemical Society*, volume 258. Amer Chemical Soc 1155 16TH ST, NW, WASHINGTON, DC 20036 USA.
- Piroi, F.; Lupu, M.; Hanbury, A.; Sexton, A. P.; Magdy, W.; and Filippov, I. V. 2010. CLEF-IP 2010: Retrieval Experiments in the Intellectual Property Domain. In *CLEF (notebook papers/labs/workshops)*.
- Qian, Y.; Guo, J.; Tu, Z.; Li, Z.; Coley, C. W.; and Barzilay, R. 2023. MolScribe: Robust Molecular Structure Recognition with Image-to-Graph Generation. *Journal of Chemical Information and Modeling*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Rajan, K.; Brinkhaus, H. O.; Agea, M. I.; Zielesny, A.; and Steinbeck, C. 2023. DECIMER. ai: an open platform for automated optical chemical structure identification, segmentation and recognition in scientific publications. *Nature communications*, 14(1): 5045.
- Sadawi, N. M.; Sexton, A. P.; and Sorge, V. 2012. Chemical structure recognition: a rule-based approach. In *Document recognition and retrieval XIX*, volume 8297, 101–109. SPIE.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shazeer, N. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Smolov, V.; Zentsev, F.; and Rybalkin, M. 2011. Imago: Open-Source Toolkit for 2D Chemical Structure Image Recognition. In *TREC*, volume 2, 6.

Staker, J.; Marshall, K.; Abel, R.; and McQuaw, C. M. 2019. Molecular structure extraction from documents using deep learning. *Journal of chemical information and modeling*, 59(3): 1017–1029.

Stärk, H.; Beaini, D.; Corso, G.; Tossou, P.; Dallago, C.; Günnemann, S.; and Liò, P. 2022. 3d infomax improves gnn for molecular property prediction. In *International Conference on Machine Learning*, 20479–20502. PMLR.

Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.

Valko, A. T.; and Johnson, A. P. 2009. CLiDE Pro: the latest generation of CLiDE, a tool for optical chemical structure recognition. *Journal of chemical information and modeling*, 49(4): 780–787.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, B.; Xu, C.; Zhao, X.; Ouyang, L.; Wu, F.; Zhao, Z.; Xu, R.; Liu, K.; Qu, Y.; Shang, F.; Zhang, B.; Wei, L.; Sui, Z.; Li, W.; Shi, B.; Qiao, Y.; Lin, D.; and He, C. 2024. MinerU: An Open-Source Solution for Precise Document Content Extraction. arXiv:2409.18839.

Wang, Y.; Wang, J.; Cao, Z.; and Barati Farimani, A. 2022. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 1–9.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.

Weininger, D. 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.

Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; and Pande, V. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science*, 9(2): 513–530.

Xia, J.; Zhao, C.; Hu, B.; Gao, Z.; Tan, C.; Liu, Y.; Li, S.; and Li, S. Z. 2023. Mole-BERT: Rethinking Pre-training Graph Neural Networks for Molecules. In *The Eleventh International Conference on Learning Representations*.

Yu, Q.; Zhang, Y.; Ni, Y.; Feng, S.; Lan, Y.; Zhou, H.; and Liu, J. 2024. Multimodal Molecular Pretraining via Modality Blending. In *The Twelfth International Conference on Learning Representations*.

Zhang, B.; and Sennrich, R. 2019. Root mean square layer normalization. *Advances in neural information processing systems*, 32.