Toward Synthesizability-Aware Multi-Step Retrosynthetic Planning

Yujie Chen[1,2], **AJie Lin**[1,2], **Tengfei Ma**[1,2], **Shu Wu**[3], **Leyi Wei**[4], **Yiping Liu**[1,2,5,*] and **Xiangxiang Zeng**[1,2,5]

[1]State Key Lab. of Chemo & Biosensing, Coll. of Comp. Sci. & Electron. Eng., Hunan University, China
[2]MOE Key Lab. of Fusion Comput. of Supercomput. & Artif. Intell., Hunan University, China
[3]NLPR, MAIS, Institute of Automation, Chinese Academy of Sciences, China
[4]Centre for AIDD, Faculty of Applied Science, Macao Polytechnic University, China
[5]Yuelushan Center for Industrial Innovation, Changsha, China
yjchen@hnu.edu.cn, yiping0liu@gmail.com

Abstract

Multi-step retrosynthetic planning aims to decompose target molecules into available starting materials by iteratively invoking single-step prediction models within a search algorithm. The success of retrosynthetic planning depends on the joint guidance of single-step reasoning and global search across steps. However, most existing frameworks make step-wise decisions based only on the current molecular state, without explicitly modeling synthesizability signals that reflect long-range reachability. In this work, we propose **GuideRetro**, a synthesizability-aware framework for multi-step retrosynthetic planning that integrates global synthesizability knowledge into step-wise retrosynthetic prediction. GuideRetro learns transferable knowledge from large-scale reaction networks by modeling the evolution of synthetic complexity along reaction pathways. During planning, a route-aware synthesis state modeling module combines the evolving retrosynthetic route with retrieved global signals to guide reactant generation at each step. Experiments on benchmark datasets show that GuideRetro achieves state-of-the-art performance. The integration of global synthesizability knowledge and route-aware modeling improves planning accuracy and search efficiency under realistic retrosynthetic settings. The code is available at: https://github.com/L-AJ/GuideRetro.

1 Introduction

Retrosynthetic planning identifies feasible synthetic routes that iteratively decompose target molecules into accessible starting materials [Corey and Wipke, 1969; Corey, 1991], forming the foundation of organic synthesis across applications in drug discovery, materials science, and industrial chemistry [Coley *et al.*, 2018a; Genheden *et al.*, 2020; Dong *et al.*, 2022]. With the growing availability of large-scale reaction datasets extracted from patents and the chemical literature [Lowe, 2012], data-driven retrosynthesis methods have emerged as the dominant paradigm in computer-aided synthesis planning (CASP).

Most data-driven multi-step planning methods formulate the task as an iterative search process, as illustrated in Figure 1a. A single-step model predicts plausible precursor sets for a given molecule at each step [Segler and Waller, 2017; Coley *et al.*, 2017; Dai *et al.*, 2019; Karpov *et al.*, 2019; Liu *et al.*, 2023b; Han *et al.*, 2024; Xu *et al.*, 2025], while a search algorithm repeatedly applies this model to expand intermediate states until purchasable starting materials are reached [Segler *et al.*, 2018; Chen *et al.*, 2020; Genheden *et al.*, 2020; Xie *et al.*, 2022; Liu *et al.*, 2023a; Wang and Montana, 2025].

Within this framework, single-step prediction defines the reaction space to be explored, while the search algorithm determines which intermediate candidates are expanded during planning. Despite progress in both components, a fundamental mismatch remains between single-step training objectives and multi-step planning goals. Existing retrosynthesis models are typically optimized to maximize the likelihood of locally plausible reactions, while multi-step planning requires selecting transformations that lead to reachable starting materials over multiple steps [Torren-Peraire *et al.*, 2024]. This misalignment implies that high confidence at the single-step level does not translate into successful multi-step planning, especially when early decisions constrain subsequent planning steps. This motivates us to design a synthesizability-aware retrosynthesis framework that explicitly incorporates global synthesizability signals to guide single-step predictions toward effective multi-step planning (Figure 1b).

To this end, we propose **GuideRetro**, a synthesizability-aware framework for multi-step retrosynthetic planning. GuideRetro learns transferable *global synthesizability knowledge* from large-scale reaction networks via preference-based optimization, providing directional signals that guide transformations leading to reachable starting materials. Given the path-dependent nature of retrosynthetic feasibility, GuideRetro further conditions this global guidance on the evolving reaction route through a *Route-aware Synthesis State Modeling* module, producing a synthesis-aware state that integrates molecular structure with route context. This

*Corresponding author.

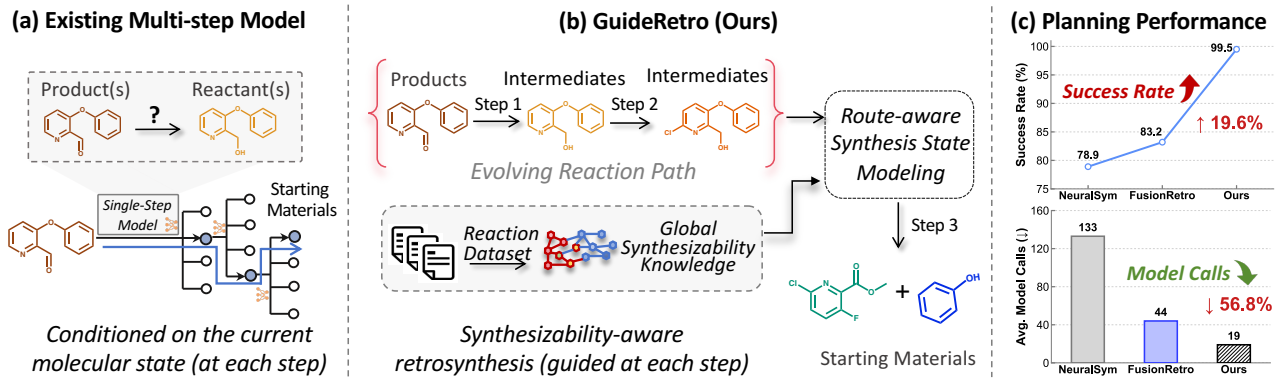

Figure 1: **(a–b)** Overview of multi-step retrosynthetic planning paradigms. **(c)** GuideRetro achieves higher planning success rates and better search efficiency than existing methods in (a).

state is then used to guide step-wise reactant generation via a *Reactants Generator*, enabling locally accurate predictions that aligned with route-level feasibility. Experiments on benchmark datasets show that GuideRetro achieves state-of-the-art multi-step planning performance, with higher planning accuracy and lower search cost, as illustrated in Figure 1c. Our main contributions are summarized as follows:

- We propose a novel synthesizability-aware retrosynthesis framework that integrates global synthesis knowledge into step-wise prediction for multi-step planning.
- We introduce a global synthesizability representation learned via preference supervision on large-scale reaction networks, capturing route-level preferences over retrosynthetic transformations.
- Extensive experiments show that GuideRetro achieves state-of-the-art performance on benchmark datasets, improving planning accuracy, robustness to route depth, and success rate under constrained search budgets.

2 Related Work

Single-step Retrosynthesis Models. Single-step retrosynthesis focuses on predicting plausible reactant sets for a given target molecule and serves as a fundamental component of data-driven synthesis planning. Early approaches rely on reaction templates to ensure chemical validity [Segler and Waller, 2017; Coley *et al.*, 2017; Dai *et al.*, 2019; Chen and Jung, 2021; Xie *et al.*, 2023], while semi-template-based methods decompose the task by first predicting bond disconnections and then completing synthons to recover full reactant structures [Shi *et al.*, 2020; Somnath *et al.*, 2021; Wang *et al.*, 2021; Wang *et al.*, 2023]. More recently, template-free models formulate retrosynthesis as a sequence-to-sequence problem and leverage Transformer architectures to directly translate product into reactants [Schwaller *et al.*, 2019; Zhong *et al.*, 2023; Han *et al.*, 2024; Wang and Montana, 2025]. Although these models have achieved strong performance at the single-step level, they are typically optimized for local reaction plausibility and do not explicitly consider whether the predicted intermediates remain feasible when integrated into multi-step retrosynthetic routes.

Search Algorithms in Retrosynthetic Planning. Multi-step retrosynthetic planning is commonly formulated as a structured search problem, in which a single-step retrosynthesis model proposes expansions and a planner selects which intermediate states to explore. Representative paradigms include Monte Carlo Tree Search (MCTS) [Segler *et al.*, 2018], proof-number search over AND/OR trees [Kishimoto *et al.*, 2019], and A*-like planners such as Retro*, which select expansions guided by learned value estimates [Chen *et al.*, 2020]. Subsequent work has focused on improving search efficiency and reliability by enhancing value estimation, for example through graph-based aggregation of partial synthesis trees in Retro*-style planning [Xie *et al.*, 2022]. More recently, learning-based strategies have been explored to refine planning behavior, including reinforcement learning for search control [Yu *et al.*, 2022], bidirectional search schemes [Yu *et al.*, 2024], and preference-based objectives learned from successful planning trajectories [Liu *et al.*, 2024; Wang and Montana, 2025]. Despite these advances, most approaches treat expansion and selection as largely decoupled, with route-level information used mainly to rank candidate expansions during search, while single-step generation remains driven by local reaction plausibility. These limitations motivate the development of synthesizability-aware methods that can align step-wise retrosynthetic predictions with route-level feasibility.

3 Preliminaries

Problem Formulation. We formulate retrosynthetic planning as a sequential decision-making problem that decomposes a target molecule into commercially available starting materials. The space of all molecules is denoted by $\mathcal{M}$, and commercially available starting materials by $\mathcal{S} \subset \mathcal{M}$. Given a target product molecule $m_0 \in \mathcal{M}$, the task is to find a retrosynthetic route whose leaf nodes belong to $\mathcal{S}$. At step $t \geq 0$, the route up to step t is represented as an ordered sequence $\tau_t = (m_0, m_1, \ldots, m_t)$, where each $m_i \in \mathcal{M}$ is an intermediate molecule obtained by applying a single-step retrosynthetic reaction to its predecessor m_{i-1}. A planning algorithm iteratively expands intermediate molecules until all terminal molecules in the route belong to $\mathcal{S}$.

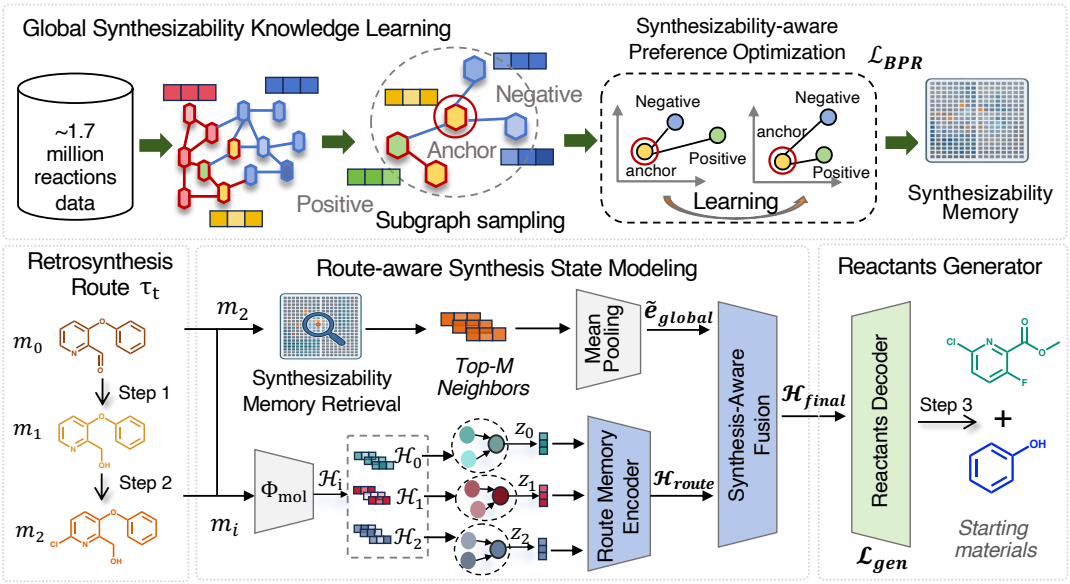

Figure 2: Overview of the GuideRetro framework. **(Top)** Synthesizability Memory is learned from large-scale reaction data via synthesizability-aware preference optimization. **(Bottom)** Given a partial route $\tau_t = (m_0, \ldots, m_t)$, GuideRetro retrieves global synthesizability guidance for m_t and fuses it with route context $\mathcal{H}_{\text{route}}(\tau_t)$ to form a synthesis-aware state $\mathcal{H}_{\text{final}}$ for reactant generation.

Synthetic Complexity Score. Synthetic Complexity Score (SCScore) [Coley *et al.*, 2018b] is a learned metric that maps a molecule m to a scalar score $s(m) \in [1, 5]$, designed to correlate with the expected number of reaction steps required for synthesis. It is trained on large-scale reaction data under the assumption that reaction products are, on average, more synthetically complex than their reactants. In this work, we use SCScore as a supervisory signal to capture global synthesizability information over the reaction network. Rather than using SCScore solely as a local heuristic, we leverage it to model how synthetic difficulty evolves along multi-step retrosynthetic pathways, thereby providing route-level guidance beyond single-step reaction feasibility.

4 The Proposed GuideRetro

This section introduces GuideRetro, a synthesizability-aware framework for multi-step retrosynthetic planning. As shown in Figure 2, GuideRetro first learns **global synthesizability knowledge** by training a reaction graph encoder with synthesizability-aware preference supervision, capturing transformation patterns that reflect long-range synthetic feasibility. During planning, a **route-aware synthesis state modeling** module combines the evolving retrosynthetic route with retrieved global guidance to construct synthesis-aware states. A **Reactant Generator** finally performs autoregressive reactant generation conditioned on these synthesis-aware states. Each component is described in detail below.

4.1 Global Synthesizability Knowledge Learning

Effective multi-step planning benefits from global and directional signals of molecular synthesizability that cannot be obtained from single-step reaction prediction alone. To obtain such signals, we learn global synthesizability knowledge from large-scale reaction data by constructing a reaction graph and learning molecule representations that model the evolution of synthetic difficulty along reaction pathways.

Synthesizability-Aware Reaction Graph Construction. A reaction graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is constructed from a reaction dataset, where the node set $\mathcal{V}$ corresponds to molecules in the molecular space $\mathcal{M}$. For each reaction, a directed edge $(u, v) \in \mathcal{E}$ is added from each reactant molecule u to the corresponding product molecule v, representing a forward reaction direction. To encode directional synthetic difficulty, each edge (u, v) is annotated with a synthetic complexity difference, which is defined as follows:

$$\Delta s_{uv} = s(v) - s(u), \tag{1}$$

where $s(\cdot)$ denotes the synthetic complexity score, and larger values of Δs_{uv} indicate more favorable retrosynthetic disconnections, as they correspond to greater reductions in molecular complexity from product v to reactant u.

Synthesizability Representation Learning. Given the reaction graph $\mathcal{G}$, a global synthesizability embedding space is learned to capture how synthetic complexity evolves along reaction pathways. To encode the inherent directionality of chemical transformations, molecular representations are initialized with a translational embedding [Bordes *et al.*, 2013], which imposes an asymmetric structure between reactants and products. Building on this directional initialization, we apply a relational graph convolutional network (RGCN) [Schlichtkrull *et al.*, 2018] over $\mathcal{G}$ to propagate synthesizability signals across multi-step synthetic pathways. At layer l, the representation of molecule i is updated by aggregating information from its incoming reactant neighbors, which is defined as:

$$\tilde{\mathbf{e}}_i^{(l+1)} = \phi^{(l)}\left(\sum_{j \in \mathcal{N}_i} c_{ij}\, W^{(l)} \tilde{\mathbf{e}}_j^{(l)} \;+\; W_0^{(l)} \tilde{\mathbf{e}}_i^{(l)}\right), \tag{2}$$

where $\tilde{\mathbf{e}}_i^{(l)} \in \mathbb{R}^{d_g}$ denotes the global synthesizability embedding, $\mathcal{N}_i = \{ j \mid (j,i) \in \mathcal{E} \}$ denotes the set of predecessor molecules that can produce molecule i, $W^{(l)}$ and $W_0^{(l)}$ are learnable transformations, c_{ij} is a degree-based normalization coefficient, and $\phi^{(l)}(\cdot)$ is a non-linear activation function.

Synthesizability-aware Preference Optimization. To align the learned representations with retrosynthetic preferences, we introduce a training objective that encourages higher scores for reaction steps leading to simpler precursors. This is based on the assumption that larger reductions in molecular complexity correspond to more effective retrosynthetic disconnections. For each product molecule v in a sampled subgraph, the reactant u^+ with the largest synthetic complexity reduction Δs_{uv} is treated as the positive sample. K negative reactants $\{u_k^-\}_{k=1}^K$ are sampled from the remaining incoming edges of v within the same mini-batch. Each candidate reaction (u, v) is scored using a bilinear interaction function [Yang *et al.*, 2014], defined as $f(u,v) = \langle \tilde{\mathbf{e}}_u, \mathbf{w}, \tilde{\mathbf{e}}_v \rangle$, where $\tilde{\mathbf{e}}_u$ and $\tilde{\mathbf{e}}_v$ are the learned representations of the reactant and product molecules, and $\mathbf{w}$ is a trainable relation vector. The model is optimized with a Bayesian Personalized Ranking loss [Rendle *et al.*, 2012]:

$$\mathcal{L}_{\mathrm{BPR}} = -\mathbb{E}\left[\frac{1}{K}\sum_{k=1}^{K} \log \sigma\left(f(u^+, v) - f(u_k^-, v)\right)\right], \quad (3)$$

which encourages the model to assign higher scores to transitions associated with greater reductions in complexity. This learning process results in an embedding space that encodes global preferences over synthesizable retrosynthetic transformations, which is used to support subsequent planning.

4.2 Route-Aware Synthesis State Modeling

To model route-dependent transformations in multi-step retrosynthesis, we introduce a route-aware synthesis state modeling module that constructs a planning state by combining route context and global synthesizability signals to guide step-wise reactant generation.

Route Memory Encoder. At retrosynthetic step t, the current molecule is m_t with partial route $\tau_t = (m_0, \ldots, m_t)$. Each molecule along the route is encoded into token-level representations using a shared molecular encoder Φ_{mol}:

$$\mathbf{H}_i = \Phi_{\mathrm{mol}}(m_i), \quad \mathbf{H}_i \in \mathbb{R}^{L_i \times d}. \quad (4)$$

To model route-level dependencies, token representations are summarized into molecule embeddings via mean pooling, $\mathbf{z}_i = \mathrm{Pool}(\mathbf{H}_i) \in \mathbb{R}^d$. Given the current embedding $\mathbf{z}_t$ and the route memory $\{\mathbf{z}_i\}_{i=0}^t$, the route-aware context representation is computed as follows:

$$\alpha_{t,i} = \frac{\exp(\mathbf{z}_t^\top \mathbf{z}_i)}{\sum_{j=0}^{t} \exp(\mathbf{z}_t^\top \mathbf{z}_j)}, \qquad \tilde{\mathbf{z}}_t = \sum_{i=0}^{t} \alpha_{t,i} \mathbf{z}_i. \quad (5)$$

The context vector is then linearly projected and fused with the token-level representation of the current molecule:

$$\mathbf{H}_{\mathrm{route}} = \mathbf{H}_t + \tilde{\mathbf{z}}_t W_r, \quad (6)$$

where $W_r \in \mathbb{R}^{d \times d}$ is a learnable projection matrix whose output is broadcast along the sequence dimension. The resulting $\mathbf{H}_{\mathrm{route}}$ serves as the route-aware state for subsequent global synthesizability fusion and reactant generation.

Global Synthesizability Retrieval. The Synthesizability Memory learned from the reaction graph stores global representations that capture synthetic feasibility patterns across chemical space. Given a query molecule m_t, we retrieve its top-M nearest neighbors based on Tanimoto similarity over Morgan fingerprints [Rogers and Hahn, 2010].The retrieved embeddings are given by:

$$\{\tilde{\mathbf{e}}_i\}_{i=1}^M = \mathrm{Retrieve}(m_t; \tilde{\mathbf{E}}), \quad (7)$$

where $\tilde{\mathbf{E}}$ denotes the Synthesizability Memory and $\tilde{\mathbf{e}}_i \in \mathbb{R}^{d_g}$ are the retrieved molecule embeddings, The retrieved embeddings are then aggregated by mean pooling, $\tilde{\mathbf{e}}_{\mathrm{global}} = \frac{1}{M}\sum_{i=1}^M \tilde{\mathbf{e}}_i$, to obtain a global synthesizability representation for the current planning step.

Synthesis-Aware Fusion. To inject global synthesizability guidance into step-wise retrosynthetic prediction, we fuse the route-contextualized token representations with the retrieved global embedding using a gated fusion mechanism. The route-contextualized token representations are denoted by $\mathbf{H}_{\mathrm{route}} \in \mathbb{R}^{L \times d}$, while the pooled global synthesizability embedding is denoted as $\tilde{\mathbf{e}}_{\mathrm{global}} \in \mathbb{R}^{d_g}$. The global embedding is then broadcast to the token level and used to compute a token-wise fusion gate:

$$\mathbf{G} = \sigma(\mathbf{H}_{\mathrm{route}} W_g + \tilde{\mathbf{e}}_{\mathrm{global}} U_g), \quad (8)$$

where σ is the sigmoid function and W_g, U_g are learnable projection matrices. The final synthesis-aware representation is then computed as:

$$\mathbf{H}_{\mathrm{final}} = \mathbf{H}_{\mathrm{route}} + \mathbf{G} \odot \phi(\tilde{\mathbf{e}}_{\mathrm{global}} W_f), \quad (9)$$

where $\phi(\cdot)$ denotes a non-linear activation and W_f is a learnable transformation matrix, and $\odot$ denotes element-wise multiplication. The resulting $\mathbf{H}_{\mathrm{final}}$ adaptively integrates route context with global synthesizability guidance and conditions the decoder for next-step reactant generation.

4.3 Reactants Generator

Given the synthesis-aware representation $\mathbf{H}_{\mathrm{final}}$, an autoregressive Transformer decoder is used to generate reactant SMILES [Weininger, 1988]. The decoder attends to $\mathbf{H}_{\mathrm{final}}$ via cross-attention, allowing each decoding step to incorporate route context and global synthesizability guidance.This design enables the model to generate chemically valid precursor sequences that follow the evolving retrosynthetic route.

4.4 Training and Inference

Training of GuideRetro is performed in two stages. We first train the Synthesizability Memory on the reaction graph using the synthesizability-aware preference objective in Eq. (3). After pretraining, the memory is fixed and used as a retrieval module during downstream retrosynthetic modeling. The remaining components are then jointly optimized with teacher forcing by minimizing the negative log-likelihood of the target reactant sequence, which is defined as:

$$\mathcal{L}_{\mathrm{gen}} = -\sum_{t=1}^{L} \log P(y_t \mid y_{<t}, \mathbf{H}_{\mathrm{final}}; \Theta), \quad (10)$$

Algorithm 1 Multi-step retrosynthetic inference

Input: Target molecule m_0, Starting Material Set $\mathcal{S}$, Synthesizability Memory $\tilde{\mathbf{E}}$
Output: Predicted reactant set $\mathcal{R}$

Initialize reactant set $\mathcal{R} \leftarrow \emptyset$.
Initialize route set $\mathcal{L} \leftarrow \{\tau_0 = (m_0)\}$.
while $\mathcal{L} \neq \emptyset$ **do**
  Select a partial route $\tau_t = (m_0, \ldots, m_t)$ from $\mathcal{L}$.
  // Step 1: Synthesizability Memory Retrieval
  Retrieve global synthesizability embedding $\tilde{\mathbf{e}}_{\text{global}}$ for current molecule m_t from $\tilde{\mathbf{E}}$.
  // Step 2: Route-aware Synthesis State Modeling
  Compute $\mathbf{H}_{\text{final}}$ from τ_t and $\tilde{\mathbf{e}}_{\text{global}}$.
  // Step 3: Reactant Generation
  Generate reactant candidates R_{pred} from $\mathbf{H}_{\text{final}}$.
  for each reactant $r \in R_{\text{pred}}$ **do**
    if $r \in \mathcal{S}$ **then**
      Add r to $\mathcal{R}$.
    else
      Extend route $\tau' \leftarrow (\tau_t, r)$.
      Add τ' to $\mathcal{L}$.
    end if
  end for
end while
return $\mathcal{R}$

where L is the length of the reactant sequence, y_t denotes the t-th token in the reactant sequence and Θ denotes the trainable parameters. During inference, GuideRetro is integrated into a multi-step planner, generating top-N precursor candidates at each expansion step until all terminal molecules are commercially available (Algorithm 1).

5 Experiments

We explore the following research questions: (1) **RQ1:** Does GuideRetro outperform baselines in multi-step retrosynthetic planning? (2) **RQ2:** How does each component affect the performance of GuideRetro? (3) **RQ3:** How well does GuideRetro generalize across diverse datasets and settings?

5.1 Experimental Setup

Datasets. The reaction graph G is constructed from the USPTO-FULL dataset [Dai *et al.*, 2019], after removing reactions involving molecules in the benchmark test sets to prevent information leakage. GuideRetro is evaluated on **RetroBench** [Liu *et al.*, 2023b], which contains $46{,}458/5{,}803/5{,}838$ training/validation/test targets. Starting materials are extracted from RetroBench as the in-stock set, and set-wise exact match accuracy is reported. To assess out-of-distribution generalization, we further evaluate on three benchmark datasets: **Retro*-190** [Chen *et al.*, 2020], **ChEMBL-1000** [Liu *et al.*, 2023a], and **GDB17-1000** [Liu *et al.*, 2023a]. Following the Retro* settings, we use commercially available building blocks from eMolecules[1] as the in-stock inventory and report search success rate. Further details are provided in Appendix A.1 and A.3.

[1]eMolecules:http://downloads.emolecules.com/free/.

Evaluation Metrics. We evaluate planning performance using standard metrics from prior work[Liu *et al.*, 2023b; Chen *et al.*, 2020]. On RetroBench, we report set-wise Top-k exact match accuracy, where a prediction is correct if the predicted starting-material set matches any reference route. On Retro*-190, ChEMBL-1000, and GDB17-1000, we report search success rate, defined as the fraction of targets for which at least one complete route whose leaf molecules are commercially available is found under a given search budget.

Baselines. We compare GuideRetro with a diverse set of representative baselines, including template-based methods (RetroSim [Coley *et al.*, 2017], NeuralSym [Segler and Waller, 2017], GLN [Dai *et al.*, 2019]), semi-template-based approaches (G2Gs [Shi *et al.*, 2020], GraphRetro [Somnath *et al.*, 2021]), and template-free models (Transformer [Karpov *et al.*, 2019], MEGAN [Sacha *et al.*, 2021], FusionRetro [Liu *et al.*, 2023b], and FusionRetro with CREAM [Liu *et al.*, 2024]). For multi-step planning, these single-step models are integrated with three established search algorithms: Retro* [Chen *et al.*, 2020], Retro-0* (i.e., Retro* without learned value estimation), and Greedy DFS. All methods are evaluated under the same settings for fair comparison.

Implementation Details. GuideRetro is trained using the Adam optimizer with a learning rate of 1×10^{-4} and batch size 32 under teacher forcing. The molecular encoder and reactant generator adopt a 3-layer Transformer architecture with embedding dimension $d = 64$, 10 attention heads, feed-forward hidden size 512, and dropout rate 0.1, with parameters shared across route steps. The Synthesizability Memory is learned by training a two-layer RGCN with embedding dimension $d_g = 512$ on the reaction graph using the synthesizability-aware preference objective in Eq. (3). After pretraining, the Synthesizability Memory is fixed and indexed using FAISS [Douze *et al.*, 2025] for efficient nearest-neighbor retrieval during downstream planning. All experiments are conducted on a single NVIDIA RTX 4090 GPU (24GB). Additional details are provided in the Appendix A.4.

5.2 Main Results (RQ1)

Route Planning Accuracy. We compare GuideRetro with representative methods on RetroBench under different search algorithms (Table 1). GuideRetro achieves the best overall performance across evaluated search settings. Under Retro*, GuideRetro reaches a Top-1 accuracy of 42.9%, improving over FusionRetro with CREBM (39.4%) by 8.9% relative gain, and maintains improvements at higher Top-k. In addition, GuideRetro reduces the performance gap to template-based methods (e.g., NeuralSym) without relying on predefined reaction templates, and achieves the highest Top-1 accuracy under Greedy DFS. These results indicate that incorporating global synthesizability knowledge with route-aware state modeling improves multi-step retrosynthetic planning.

Planning Robustness under Increasing Route Depth. Beyond overall planning accuracy, we further analyze the impact of route depth on planning success rates using Greedy DFS. As shown in Figure 4, GuideRetro consistently achieves higher success rates than all template-free baselines across

Single-step Models \ Search Algorithm	Retro*					Retro*-0					Greedy DFS
	Top-1	Top-2	Top-3	Top-4	Top-5	Top-1	Top-2	Top-3	Top-4	Top-5	Top-1
Template-based											
RetroSim [Coley *et al.*, 2017]	35.1	40.5	42.9	44.0	44.6	35.0	40.5	43.0	44.1	44.6	31.5
NeuralSym [Segler and Waller, 2017]	41.7	49.2	52.1	53.6	54.4	42.0	49.3	52.0	53.6	54.3	39.2
GLN [Dai *et al.*, 2019]	39.6	48.9	52.7	54.6	55.7	39.5	48.7	52.6	54.5	55.6	38.0
Semi-template-based											
G2Gs [Shi *et al.*, 2020]	5.4	8.3	9.9	10.9	11.7	4.2	6.5	7.6	8.3	8.9	3.8
GraphRetro [Somnath *et al.*, 2021]	15.3	19.5	21.0	21.9	22.4	15.3	19.5	21.0	21.9	22.2	14.4
Template-free											
Transformer [Karpov *et al.*, 2019]	31.3	40.4	44.7	47.2	48.9	31.2	40.5	45.1	47.3	48.7	26.7
MEGAN [Sacha *et al.*, 2021]	18.8	27.9	32.7	36.6	38.1	18.6	27.7	32.6	36.4	38.5	32.9
FusionRetro [Liu *et al.*, 2023b]	37.5	45.0	48.2	50.0	50.9	37.5	45.0	48.3	50.2	51.2	33.8
FusionRetro + CREBM [Liu *et al.*, 2024]	39.4	46.6	49.3	50.7	51.5	39.6	46.7	49.5	51.0	51.7	-
RetroInText	41.2	48.7	51.8	53.3	54.2	42.1	49.9	53.0	54.7	55.7	39.8
GuideRetro (Ours)	**42.9**	**50.7**	**54.3**	**56.0**	**57.1**	**42.9**	**50.7**	**54.3**	**55.9**	**57.0**	**40.9**

Table 1: Retrosynthetic planning results on RetroBench dataset with Set-wise Exact Match Accuracy(%) *(Best in **bold**, second underlined)*.

Single-step	Search	Retro*-190				ChEMBL-1000				GDB17-1000			
		50	100	200	500	50	100	200	500	50	100	200	500
NeuralSym	Retro*-0	26.8	37.4	58.9	78.9	66.5	69.1	72.2	75.0	4.0	5.2	6.5	7.6
	Retro*	43.7	51.6	66.3	85.8	67.2	71.0	72.9	76.3	3.6	5.7	6.8	8.3
FusionRetro	Retro*-0	66.3	72.6	77.4	83.2	70.6	78.1	82.8	86.1	77.6	86.2	91.6	95.0
	Retro*	85.3	87.9	88.9	90.0	84.9	86.6	87.8	88.1	90.6	93.1	95.1	96.1
GuideRetro (Ours)	Retro*-0	91.1	97.9	99.5	99.5	93.3	96.2	96.8	96.8	97.7	99.5	99.6	99.6
	Retro*	**99.5**	**100.0**	**100.0**	**100.0**	**96.3**	**96.6**	**96.8**	**96.8**	**99.2**	**99.5**	**99.6**	**99.6**

Table 2: Search Success Rates (%) on Retro*-190, ChEMBL-1000, and GDB17-1000 datasets with varying search budgets.

different route depths. Although the performance of all methods degrades as route depth increases due to error propagation across successive steps, GuideRetro maintains a clear performance advantage at every depth. In particular, at a depth of 8, GuideRetro achieves a 29.4% success rate, nearly doubling the performance of FusionRetro (14.8%), while the vanilla Transformer and GraphRetro exhibit near-zero success rates (≈0%). These results indicate that incorporating global synthesizability knowledge into route-aware synthesis state modeling leads to more informative intermediate decisions and improved robustness in deeper retrosynthetic planning.

Planning Success on Challenging Targets. To further assess planning reliability on difficult targets, we evaluate GuideRetro on the Retro*-190, which contains 190 challenging molecules. As shown in Table 2, GuideRetro achieves higher success rates than prior methods under limited search budgets, improving Retro*-0 success from 66.3% to 91.1% at budget 50 and reaching 99.5% under Retro* at the same budget. These results validate the effectiveness of GuideRetro for reliable multi-step retrosynthetic planning.

5.3 Ablation Study (RQ2)

To evaluate the contribution of each component in GuideRetro, we perform ablation studies on RetroBench using Greedy DFS (Table 3). Removing the Route Memory Encoder (**w/o RME**) leads to the largest drop in exact match accuracy from 40.9% to 34.0%, underscoring the necessity of modeling the evolving retrosynthetic route. Excluding Global Synthesizability Retrieval (**w/o GSR**) reduces accuracy to 33.8%, showing that route context alone is insufficient without global guidance. Removing the synthesizability-aware preference objective (**w/o SPO**) causes a smaller but consistent decline to 36.6%, indicating the benefit of preference-aligned pretraining. Although the gain from GSR is moderate, Appendix B.2 shows that its impact increases with route depth, indicating complementary effects across components.

5.4 Generalization Ability of GuideRetro (RQ3)

To evaluate the generalization of GuideRetro, we integrate it with Retro* and Retro*-0 search on two out-of-distribution benchmarks: ChEMBL-1000 and GDB17-1000. As shown in Table 2, GuideRetro consistently achieves higher search success rates than prior methods, with the largest improvements observed under constrained search budgets. For example, under Retro*-0 with a budget of 50, GuideRetro improves success rates from 70.6% to 93.3% on ChEMBL-1000 and reaches 97.0% on the more challenging GDB17-1000 benchmark, compared to 77.6% obtained by FusionRetro.

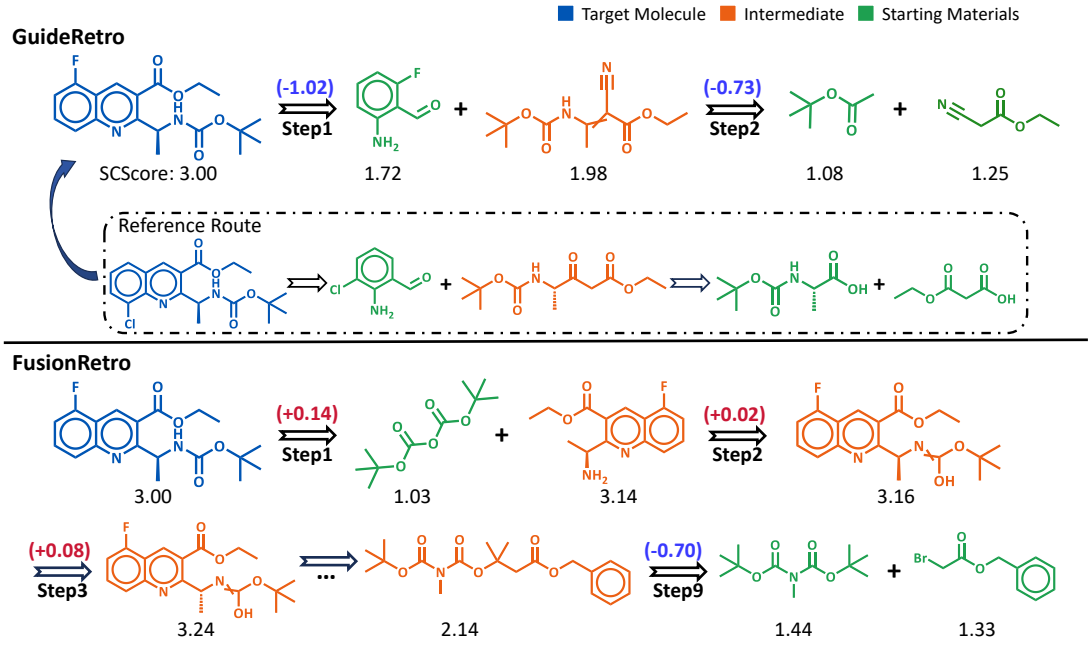

Figure 3: Retrosynthetic routes for the same Retro*-190 target molecule (SCScore = 3.00), predicted by GuideRetro (top) and FusionRetro (bottom). Arrows denote retrosynthetic steps with SCScore changes; the dashed box shows a reference route for a similar molecule.

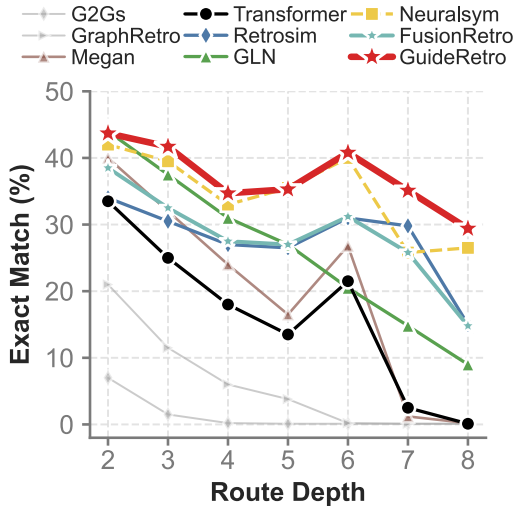

Figure 4: Planning success rate vs. route depth under Greedy DFS.

Under Retro* search, GuideRetro achieves 97.7% success on GDB17-1000 at the same budget and consistently outperforms FusionRetro across datasets and budgets. These results demonstrate that GuideRetro generalizes well across diverse chemical spaces and search settings.

5.5 Case Study

To provide a qualitative comparison, we analyze retrosynthetic routes generated by GuideRetro and FusionRetro for a target molecule (SCScore = 3.0) from Retro-190*. As shown in Figure 3, GuideRetro finds a synthesis route with only two model calls, while FusionRetro requires fifteen calls.

Model Variant	Exact Match (%)	Δ Drop
GuideRetro w/o RME	34.0	−6.9
GuideRetro w/o GSR	33.8	−7.1
GuideRetro w/o SPO	36.6	−2.9
GuideRetro (Full)	**40.9**	–

Table 3: Ablation study of GuideRetro on RetroBench.

Along the GuideRetro route, intermediate complexity decreases steadily, enabling efficient convergence to starting materials. Moreover, the resulting route is consistent with the reference synthesis in the reaction graph. This example highlights the effectiveness of GuideRetro for reliable and efficient multi-step retrosynthetic planning. More examples are provided in Appendix B.4.

6 Conclusion

We introduce **GuideRetro**, a synthesizability-aware framework for multi-step retrosynthetic planning. By learning global synthesizability knowledge from large-scale reaction networks and integrating it with route-aware synthesis state modeling, GuideRetro provides effective guidance for stepwise retrosynthetic prediction throughout the planning process. This enables planning decisions to incorporate route-level feasibility beyond local chemical context. Experiments on multiple benchmarks demonstrate consistent improvements in planning accuracy, robustness, and search efficiency, highlighting the value of synthesizability-aware guidance for reliable retrosynthesis.

Acknowledgments

This work was supported by the National Science and Technology Major Project (Grant No. 2023ZD0120902), the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (JYB2025XDXM602), the National Natural Science Foundation of China (Grant Nos. 62522110, 62472152, 62425204, U22A2037, 62450002, 62432011 and 625B2067), the Hunan Provincial Natural Science Foundation of China (Grant No. 2024JJ4015), the Project of Yuelushan Center for Industrial Innovation (Grant No. 2025YCII0214), and the Science and Technology Development Fund of MACAU (no. 0133/2024/RIB2).

Contribution Statement

Yujie Chen and Ajie Lin made equal contributions to this work.

References

[Bordes *et al.*, 2013] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26, 2013.

[Chen and Jung, 2021] Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021.

[Chen *et al.*, 2020] Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: learning retrosynthetic planning with neural guided a* search. In *International conference on machine learning*, pages 1608–1616. PMLR, 2020.

[Coley *et al.*, 2017] Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Computer-assisted retrosynthesis based on molecular similarity. *ACS central science*, 3(12):1237–1245, 2017.

[Coley *et al.*, 2018a] Connor W Coley, William H Green, and Klavs F Jensen. Machine learning in computer-aided synthesis planning. *Accounts of chemical research*, 51(5):1281–1289, 2018.

[Coley *et al.*, 2018b] Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Scscore: synthetic complexity learned from a reaction corpus. *Journal of chemical information and modeling*, 58(2):252–261, 2018.

[Corey and Wipke, 1969] Elias James Corey and W Todd Wipke. Computer-assisted design of complex organic syntheses: Pathways for molecular synthesis can be devised with a computer and equipment for graphical communication. *Science*, 166(3902):178–192, 1969.

[Corey, 1991] Elias James Corey. The logic of chemical synthesis: multistep synthesis of complex carbogenic molecules (nobel lecture). *Angewandte Chemie International Edition in English*, 30(5):455–465, 1991.

[Dai *et al.*, 2019] Hanjun Dai, Chengtao Li, Connor Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. *Advances in Neural Information Processing Systems*, 32, 2019.

[Dong *et al.*, 2022] Jingxin Dong, Mingyi Zhao, Yuansheng Liu, Yansen Su, and Xiangxiang Zeng. Deep learning in retrosynthesis planning: datasets, models and tools. *Briefings in Bioinformatics*, 23(1), 2022.

[Douze *et al.*, 2025] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. *IEEE Transactions on Big Data*, 2025.

[Genheden *et al.*, 2020] Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):70, 2020.

[Han *et al.*, 2024] Yuqiang Han, Xiaoyang Xu, Chang-Yu Hsieh, Keyan Ding, Hongxia Xu, Renjun Xu, Tingjun Hou, Qiang Zhang, and Huajun Chen. Retrosynthesis prediction with an iterative string editing model. *Nature Communications*, 15(1):6404, 2024.

[Karpov *et al.*, 2019] Pavel Karpov, Guillaume Godin, and Igor V Tetko. A transformer model for retrosynthesis. In *International conference on artificial neural networks*, pages 817–830. Springer, 2019.

[Kishimoto *et al.*, 2019] Akihiro Kishimoto, Beat Buesser, Bei Chen, and Adi Botea. Depth-first proof-number search with heuristic edge cost and application to chemical synthesis planning. *Advances in Neural Information Processing Systems*, 32, 2019.

[Liu *et al.*, 2023a] Guoqing Liu, Di Xue, Shufang Xie, Yingce Xia, Austin Tripp, Krzysztof Maziarz, Marwin Segler, Tao Qin, Zongzhang Zhang, and Tie-Yan Liu. Retrosynthetic planning with dual value networks. In *International conference on machine learning*, pages 22266–22276. PMLR, 2023.

[Liu *et al.*, 2023b] Songtao Liu, Zhengkai Tu, Minkai Xu, Zuobai Zhang, Lu Lin, Rex Ying, Jian Tang, Peilin Zhao, and Dinghao Wu. Fusionretro: molecule representation fusion via in-context learning for retrosynthetic planning. In *International Conference on Machine Learning*, pages 22028–22041. PMLR, 2023.

[Liu *et al.*, 2024] Songtao Liu, Hanjun Dai, Yue Zhao, and Peng Liu. Preference optimization for molecule synthesis with conditional residual energy-based models. *arXiv preprint arXiv:2406.02066*, 2024.

[Lowe, 2012] Daniel Mark Lowe. *Extraction of chemical structures and reactions from the literature*. PhD thesis, University of Cambridge, 2012.

[Rendle *et al.*, 2012] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*, 2012.

[Rogers and Hahn, 2010] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.

[Sacha *et al.*, 2021] Mikołaj Sacha, Mikołaj Błaz, Piotr Byrski, Paweł Dabrowski-Tumanski, Mikołaj Chrominski, Rafał Loska, Paweł Włodarczyk-Pruszynski, and Stanisław Jastrzebski. Molecule edit graph attention network: modeling chemical reactions as sequences of graph edits. *Journal of Chemical Information and Modeling*, 61(7):3273–3284, 2021.

[Schlichtkrull *et al.*, 2018] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.

[Schwaller *et al.*, 2019] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS central science*, 5(9):1572–1583, 2019.

[Segler and Waller, 2017] Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry–A European Journal*, 23(25):5966–5971, 2017.

[Segler *et al.*, 2018] Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018.

[Shi *et al.*, 2020] Chence Shi, Minkai Xu, Hongyu Guo, Ming Zhang, and Jian Tang. A graph to graphs framework for retrosynthesis prediction. In *International conference on machine learning*, pages 8818–8827. PMLR, 2020.

[Somnath *et al.*, 2021] Vignesh Ram Somnath, Charlotte Bunne, Connor Coley, Andreas Krause, and Regina Barzilay. Learning graph models for retrosynthesis prediction. *Advances in Neural Information Processing Systems*, 34:9405–9415, 2021.

[Torren-Peraire *et al.*, 2024] Paula Torren-Peraire, Alan Kai Hassen, Samuel Genheden, Jonas Verhoeven, Djork-Arné Clevert, Mike Preuss, and Igor V Tetko. Models matter: the impact of single-step retrosynthesis on synthesis planning. *Digital Discovery*, 3(3):558–572, 2024.

[Wang and Montana, 2025] Mianchu Wang and Giovanni Montana. Retrosynthesis planning via worst-path policy optimisation in tree-structured mdps. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[Wang *et al.*, 2021] Xiaorui Wang, Yuquan Li, Jiezhong Qiu, Guangyong Chen, Huanxiang Liu, Benben Liao, Chang-Yu Hsieh, and Xiaojun Yao. Retroprime: A diverse, plausible and transformer-based method for single-step retrosynthesis predictions. *Chemical Engineering Journal*, 420:129845, 2021.

[Wang *et al.*, 2023] Yu Wang, Chao Pang, Yuzhe Wang, Junru Jin, Jingjie Zhang, Xiangxiang Zeng, Ran Su, Quan Zou, and Leyi Wei. Retrosynthesis prediction with an interpretable deep-learning framework based on molecular assembly tasks. *Nature Communications*, 14(1):6155, 2023.

[Weininger, 1988] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.

[Xie *et al.*, 2022] Shufang Xie, Rui Yan, Peng Han, Yingce Xia, Lijun Wu, Chenjuan Guo, Bin Yang, and Tao Qin. Retrograph: Retrosynthetic planning with graph search. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2120–2129, 2022.

[Xie *et al.*, 2023] Shufang Xie, Rui Yan, Junliang Guo, Yingce Xia, Lijun Wu, and Tao Qin. Retrosynthesis prediction with local template retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5330–5338, 2023.

[Xu *et al.*, 2025] Li-Cheng Xu, Miao-Jiong Tang, Junyi An, Fenglei Cao, and Yuan Qi. A unified pre-trained deep learning framework for cross-task reaction performance prediction and synthesis planning. *Nature Machine Intelligence*, 7(9):1561–1571, 2025.

[Yang *et al.*, 2014] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.

[Yu *et al.*, 2022] Yemin Yu, Ying Wei, Kun Kuang, Zhengxing Huang, Huaxiu Yao, and Fei Wu. Grasp: Navigating retrosynthetic planning with goal-driven policy. *Advances in Neural Information Processing Systems*, 35:10257–10268, 2022.

[Yu *et al.*, 2024] Kevin Yu, Jihye Roh, Ziang Li, Wenhao Gao, Runzhong Wang, and Connor Coley. Double-ended synthesis planning with goal-constrained bidirectional search. *Advances in Neural Information Processing Systems*, 37:112919–112949, 2024.

[Zhong *et al.*, 2023] Weihe Zhong, Ziduo Yang, and Calvin Yu-Chian Chen. Retrosynthesis prediction using an end-to-end graph generative architecture for molecular graph editing. *Nature Communications*, 14(1):3009, 2023.