

# LLM-AUGMENTED CHEMICAL SYNTHESIS AND DESIGN DECISION PROGRAMS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Retrosynthesis, the process of breaking down a target molecule into simpler precursors through a series of valid reactions, stands at the core of organic chemistry and drug development. Although recent machine learning (ML) research has advanced single-step retrosynthetic modeling and subsequent route searches, these solutions remain restricted by the extensive combinatorial space of possible pathways. Concurrently, large language models (LLMs) have exhibited remarkable chemical knowledge, hinting at their potential to tackle complex decision-making tasks in chemistry. In this work, we explore whether LLMs can successfully navigate the highly constrained, multi-step retrosynthesis planning problem. We introduce an efficient scheme for encoding reaction pathways and present a new route-level search strategy, moving beyond the conventional step-by-step reactant prediction. Through comprehensive evaluations, we show that our LLM-augmented approach excels at retrosynthesis planning and extends naturally to the broader challenge of synthesizable molecular design.

## 1 INTRODUCTION

Retrosynthesis Corey & Wipke (1969); Corey et al. (1985) concerns with breaking down a target molecular structure into a sequence of simpler or more readily available precursor structures and chemical reactions Boström et al. (2018). It is essential for many chemistry problems that require the realization of proposed molecular structures from organic synthesis to drug discovery Blakemore et al. (2018). Nevertheless, the search space for a given target is tremendous as the number of possible synthesis pathways grows exponentially with the number of reaction steps or the depth of the route tree. Consequently, efficient decision-making in retrosynthesis planning, and more broadly, in chemical design, remains a critical challenge.

Recent research has harnessed machine learning to tackle retrosynthesis by modeling reactions with a single-step model which predicts a reaction template, i.e., a reaction coded as a pattern, to synthesize the given target molecule Segler & Waller (2017); Coley et al. (2017); Liu et al. (2017), including graph neural networks Dai et al. (2019a); Chen & Jung (2021a), and subsequently reverse the template to obtain the reactants. Another branch of single-step models do not rely on the provided reaction templates and directly predict reactants Liu et al. (2017); Schwaller et al. (2020); Igashov et al.. After training the single-step models, they are further connected with a search algorithm (e.g. Monte Carlo tree search Segler et al. (2018b) or A\* search Chen et al. (2020a) to perform multi-step retrosynthetic analysis, which halts when a path to a set of predefined purchasable molecules is found.

Recent studies have shown that large language models (LLMs) implicitly encode substantial chemical knowledge, as evidenced by their remarkable performance in searching molecular structures with optimized properties Wang et al. (2024). In addition, LLMs have been leveraged for reasoning and planning problems, such as automated experimentation in chemistry M. Bran et al. (2024); Boiko et al. (2023). Despite the apparent promise, the extent to which LLMs can handle tightly constrained decision processes, such as retrosynthesis planning, remains largely unexplored. Unlike open-ended tasks like text generation, retrosynthesis imposes rigorous constraints on the sequence of actions (reaction steps). Only certain reaction templates are valid, and only commercially available or otherwise feasible precursors can be used.

In this paper, we investigate whether the knowledge embedded in LLMs can be effectively leveraged for complex sequential decision-making tasks in chemistry such as retrosynthesis planning. Crucially

and by contrast to existing LLM works for retrosynthesis Nguyen-Van et al. (2024); Yang et al. (2024), we *do not* tune the base LLM. Instead, by exploring how LLMs perform under heavy constraints, we aim to gain insights into their potential to serve as powerful decision-making engines, ultimately advancing our understanding of their capabilities in chemistry and beyond. Furthermore, we expand the scope to study the capability of LLMs in not only finding a synthesis pathway, but also simultaneously optimizing the property of the target molecule, known as synthesizable molecular design Bradshaw et al. (2019; 2020); Gottipati et al. (2020); Horwood & Noutahi (2020); Korovina et al. (2020); Gao et al. (2022; 2024); Koziarski et al. (2024); Cretu et al. (2024); Seo et al. (2024); Swanson et al. (2024). Our main contributions are as follows:

- ▷ We propose an efficient and effective way to encode the sequence of synthesis decisions: (1) a language to describe reactions that LLMs understand and (2) efficient data structures to store the exponential-growth tree-structured synthesis pathways.
- ▷ We propose a novel way of searching sequence-level decisions with a smooth reward function and partial feedback by sampling from the space of decision sequences (full multi-step synthetic pathways) rather than individual states (a single reaction step).
- ▷ Experimentally, we study both the retrosynthesis planning and synthesizable molecular design problems in this unifying paradigm of LLM-augmented reaction decision program.

## 2 PROBLEM FORMULATION

We formulate the retrosynthesis planning problem as a sequential decision making problem. At the core of this task is a **molecule set**, which contains either the molecules we aim to synthesize (target) or purchase directly (permitted commercial building blocks). We initialize the molecule set with only the **target molecule** and evolve over successive search steps until there is no molecule in the set that is non-purchasable.

At each step, we use a **backward reaction** to decompose a molecule in the set (the product) into its reactants. This involves removing the product from the molecule set and adding the corresponding reactants generated by the selected backward reaction. The process terminates when either all molecules remaining in the set are purchasable or the maximum budget of attempts is reached.

A reaction is formally defined by a **reaction template**, which specifies a structural transformation pattern in the form of a SMARTS string Daylight Chemical Information Systems. We denote the set of feasible reaction templates by  $\mathbb{T}$  and the set of purchasable compounds by  $\mathbb{C}$ . Both  $\mathbb{T}$  and  $\mathbb{C}$  are flexible and can be refined or expanded without altering the underlying framework, ensuring adaptability to various chemical spaces.

Given  $\mathbb{T}$  and  $\mathbb{C}$ , our goal is to iteratively select backward reactions that construct a valid synthetic route for the target molecule. Each molecule in the synthetic route, including intermediates, is explicitly defined through the application of reaction templates to its reactants. Compared to general molecule generation tasks, retrosynthesis planning introduces additional challenges, such as enforcing chemical reaction rules and ensuring the use of commercially available building blocks.

### 2.1 RETROSYNTHESIS PLANNING

Given a target molecule  $M_{\text{target}}$ , the objective is to identify a sequence of reactions  $\{r_1, r_2, \dots, r_n\}$  such that:

1.  $M_{\text{target}}$  can be recursively decomposed into reactants by applying reaction templates from  $\mathbb{T}$ .
2. The final set of reactants consists exclusively of molecules in  $\mathbb{C}$ .
3. Each reaction  $r_i \in \mathbb{T}$  is chemically valid and adheres to the predefined reaction rules.

At each decision step  $t$ , we select a reaction  $r_t \in \mathbb{T}$  to apply to a molecule  $M_t$  in the molecule set. This generates its reactants  $\{M_{t,1}, M_{t,2}, \dots\}$ , which are then added to the molecule set, replacing  $M_t$ . The task is completed when all terminal nodes in the synthetic pathway correspond to molecules in  $\mathbb{C}$ .

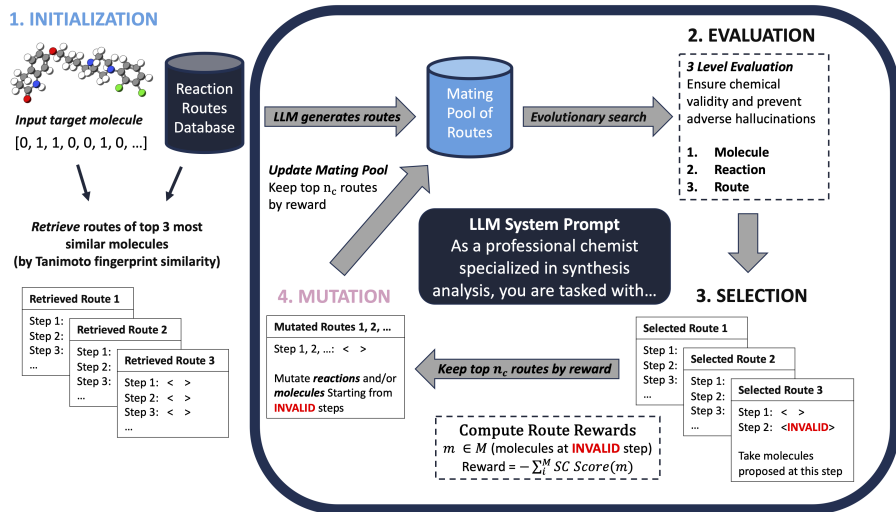


Figure 1: Overview of the LLM-Syn-Planner. **1. INITIALIZATION:** Based on the target molecule, reaction routes of similar molecules are retrieved and scored by the SC score (Coley et al., 2018). **2. EVALUATION:** The LLM generates new routes which are evaluated. **3. SELECTION:** Starting from invalid steps in the reaction routes, the SC score of the molecules at this step are computed and the top  $n_c$  routes are selected. **4. MUTATION:** Starting from these invalid steps, the LLM proposes mutations to modify the molecules and/or reactions at this step. Repeat until a solution is found or the budget is reached.

## 2.2 SYNTHESIZABLE MOLECULAR DESIGN

In contrast to retrosynthesis planning, we consider the **synthesizable molecular design** problem, where the goal is to find molecules with optimal properties evaluated by an oracle function  $O$ , while simultaneously ensuring that they are synthetically accessible through feasible reaction pathways.

$$\arg \max_{m \in \Omega} O(m) \text{ s.t. } V(R(m)) = 1$$

where  $\Omega$  is the set of generated molecules,  $R(\cdot)$  returns the synthesis path, and  $V(\cdot)$  checks the validity of the path.

## 3 METHODOLOGY

### 3.1 ROUTE FORMATTING

Traditional machine learning methods directly predict reaction classes or reactants based on input molecules, which by definition, defines the synthesis route when the full set of reaction classes and molecules are considered (Zhong et al., 2024). However, using LLMs as retrosynthesis route generators necessitates a well-defined textual input-output format, as LLMs are highly sensitive to prompt design (Sclar et al., 2024). A critical challenge lies in determining how to represent the retrosynthesis route for LLMs. Prior research has proposed two main representation formats:

- ▷ Textual descriptions (Liu et al., 2024a): Textual descriptions align naturally with the text generation capabilities of LLMs and uses descriptive language to detail each reaction step. However, the flexibility and lack of standardization in textual descriptions make it challenging to consistently extract essential information, such as reactants, products, and reactions. This ambiguity complicates the evaluation of individual steps and the validation of the overall synthesis route.
- ▷ Tree structures (Chang et al.): Tree structures (Figure 2a) represent synthetic pathways as hierarchical trees, capturing the relationships between reactants and products in a structured manner. While tree structures provide a more systematic representation, their complexity

increases significantly in multi-step retrosynthesis tasks, leading to deeply nested structures that can overwhelm the LLM’s reasoning capabilities. Additionally, the lack of an explicit reasoning process within tree representations can limit the performance of LLMs in retrosynthesis planning.

To address these limitations, we draw inspiration from traditional tree search-based approaches to retrosynthesis planning (Segler et al., 2018a). In these approaches, the nodes in the search tree represent synthetic states, and the tree itself contains all molecules required to synthesize the target molecule at the root. A target molecule is considered synthesized when all leaf nodes in the tree correspond to purchasable building blocks. The edges of the tree correspond to reactions, which specify a chemical transformation between states of connected nodes.

Building on this framework, we reformulate retrosynthesis planning into a step-by-step decision-making process that is more suitable for LLMs (Figure 2b). Specifically, we represent the synthesis route as a sequence of decisions, where each step involves proposing a reaction from a database of reaction results, i.e., *reaction templates*. The LLM maintains a dynamic molecule set that starts with the target molecule and evolves as reactions are selected, ending when all molecules in the set are purchasable. To improve the decision-making process, we integrate a reasoning component called the "Rational" at each step. This reasoning step encourages the LLM to think before making decisions (Wei et al., 2022). Additionally, we ask the LLM to explicitly output the product and reactants in each step, in order to keep the generated route more consistent.

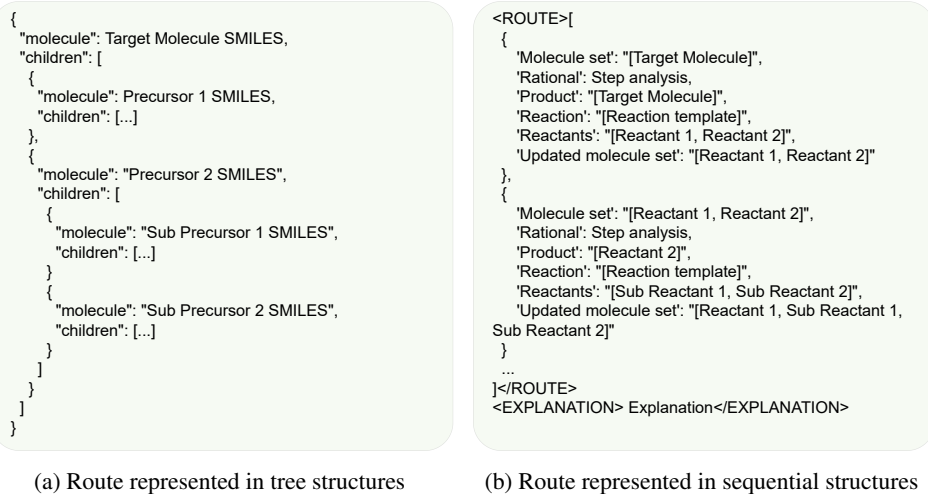


Figure 2: Different route formats of retrosynthesis routes

### 3.2 LLM AS A SINGLE-STEP PREDICTION MODEL

Recent studies have demonstrated the potential of utilizing LLMs as planners for complex decision-making tasks (Song et al., 2023; Huang et al., 2024). A common approach is integrating LLMs with traditional search algorithms such as MCTS (Zhao et al., 2024) and A\* search (Zhuang et al., 2023). This integration addresses a key limitation of LLMs: their lack of a systematic mechanism to explore structured solution spaces. Without such mechanisms, LLMs may struggle to effectively navigate complex decision-making scenarios. The core idea of these methods is straightforward: treat the LLM as a policy that directly generates the next action based on the history of past actions and observations. Meanwhile, search algorithms like MCTS and A\* systematically explore and optimize the solution space, ensuring robustness and completeness.

Building on this, we propose using LLMs as single-step retrosynthesis predictors, and operate in the template-based approach, where we start with a pre-defined templates set that represents all the reactions the LLM *could* suggest and a reference reactions database based on USPTO. The product molecule in each step serves as the input, and the LLM is queried to predict a reaction that synthesizes this product molecule. To do this, we first task the LLM with identifying substructures and functional

**Algorithm 1:** LLM-Syn-Planner Algorithm

**Data:** The target molecule  $T$ ; the reward function  $F$ ; the evaluation function  $E$ ; the population size  $n_c$ ; the number of retrieval size  $n_o$ ; the routes retrieval set  $\mathbb{O}$ ; the maximum number of attempts budget.

**Result:** Found synthesis routes population  $\mathbb{P}^*$

**begin**

$\mathbb{P}_0 = []$ ;

**while**  $len(\mathbb{P}_0) < n_c$  **do**

    sample  $\mathbb{P}_o = \{p_i\}_{i=1}^{n_o}$  from  $\mathbb{O}$  proportionally to their products' Tanimoto similarity to  $T$ ;

$\mathbb{R}_0.append(INITIALIZATION(T, \mathbb{P}_o))$ ;

**for**  $p \in \mathbb{P}_0$  **do**

    Compute  $F(p)$ ;

**for**  $t \in [1, \text{budget}]$  **do**

    offspring = [];

**for** num\_mutations **do**

      sample  $p$  from  $\mathbb{P}_t$  proportionally to reward  $F(p)$ ;

      evaluate  $p$  using the evaluation function  $E(p)$  to get feedback  $f$ ;

      offspring.append( $MUTATION(T, p, f)$ );

**for**  $p \in \mathbb{P}_t$  **do**

      Compute  $F(p)$ ;

$\mathbb{P}_{t+1} \leftarrow \text{sorted}(\mathbb{P}_t)[n_c]$ ;

  Return  $\mathbb{P}_{\text{budget}}$ ;

groups in the product molecule. Next, we draw inspiration from Coley et al. (2017) and compute the Tanimoto similarity between the substructures and the product molecules in the reference reactions database. The hypothesis is that similar product molecules are synthesized from similar reactions. Following these steps, the LLM *retrieves* a template from the pre-defined list, which is important as it removes any possibility of hallucinated templates. The template is then applied to the product molecule to obtain a set of predicted reactants.

By contrast to existing single-step prediction models Maziarz et al. (2023), it is non-trivial to obtain a probability of choosing a template from an LLM. Therefore, we assign pseudo-probabilities to the predicted reactions by employing *self-consistency frequency*, which is an ensemble approach that samples  $k$  independent reactions for the next step, denoted as  $\{r_{t+1}^{(j)}\}_{j=1}^k \sim p(r_{t+1}|m)$  at step  $t$ . From these samples, we identify the unique reactions and consider them as the set of potential next-step reactions. The frequency of each reaction in this set is then used to compute its cumulative score, given by:

$$p(n) = \frac{\#\{j \mid r_{t+1}^{(j)} = n\}}{k},$$

where  $\#\{j \mid r_{t+1}^{(j)} = n\}$  denotes the count of samples for reaction  $n$ . In essence, this expression computes how many times each reaction appears in  $k$  sampled reactions.

Finally, we integrate the LLM as a single-step predictor with MCTS (Segler et al., 2018a) or Retro\* (Chen et al., 2020b) search algorithms to explore retrosynthesis pathways.

### 3.3 LLM AS A SYNTHESIS PATHWAY SAMPLER

Although LLMs can leverage search algorithms to explore the search space, akin to existing works that pair single-step reaction prediction with search algorithms (Zhong et al., 2024), we are particularly interested in their ability to design synthesis routes directly for a given target molecule. To this end, we propose an evolutionary search algorithm named LLM-Syn-Planner that enables LLMs to *generate* and optimize the whole retrosynthetic pathways directly. We emphasize *generate* as the LLM is not explicitly retrieving a reaction template like in the case of using the LLM as a single-step prediction model and coupling a search algorithm. Unlike existing works that follow this paradigm Zhong et al. (2024), our approach generates the entire multi-step synthesis tree directly. The algorithm operates as follows: Given a target molecule, we first generate an initial pool of retrosynthetic routes using

**INITIALIZATION** queries from LLMs, where each route is evaluated using a reward function,  $F(\cdot)$ . Next, a route is sampled with a probability proportional to its reward and edited using a **MUTATION** operator to generate offspring. This mutation process is repeated  $num\_mutation$  times, after which the newly generated offspring are added to the population. The offspring are then evaluated using  $F(\cdot)$ , and the  $n_c$  fittest candidates at each step are selected to pass on to the next generation. This iterative process continues until the maximum number of model calls is reached. The overall workflow consists of four key stages: (1) **Initialization**, (2) **Evaluation**, (3) **Selection**, and (4) **Mutation**. This process is outlined in Algorithm 1.

**Initialization.** In the **INITIALIZATION** function, we query the LLM to generate initial retrosynthesis routes for the target molecule. To enhance its predictions, we employ a molecular similarity-based retrieval-augmented generation (RAG) approach, providing reference routes for the LLM. Specifically, we use the Morgan molecular fingerprint with Tanimoto similarity to identify structurally similar molecules in a database and retrieve their corresponding synthesis routes. We then provide the synthesis routes of the top three most similar molecules as references to the LLM.

**Evaluation.** We propose a three-level evaluation process to assess the quality of each step in the generated synthetic route: molecule level, reaction level, and route level, as shown in Table 1.

At the molecule level, we validate whether the molecules in the molecule set are both valid and purchasable. At the reaction level, we first perform reaction mapping to verify the reactions proposed by the LLM. This involves grounding and matching them against a reaction database. We begin by searching for exact matches. If no exact match is found, we retrieve the top 100 most similar reactions based on reaction fingerprint similarity. These candidates are then filtered by assessing whether the proposed reaction is chemically feasible for the given product molecule, as even if the retrieved route is for a similar molecule, slight differences in the target molecule structure can render the reaction incompatible. The most similar valid reaction is retained as the matched reaction. Finally, we replace the original reaction proposed by the LLM with the identified match, thus removing the possibility of a hallucinated reaction that we cannot easily verify the chemical soundness of. If no valid match is found in this process, we label the reaction as non-existent. At the route level, we evaluate route connectivity by checking whether the ‘molecule set’ in a given step aligns with the ‘updated molecule set’ from the previous step and whether the expected ‘product’ appears in the current step’s molecule set. A step is considered valid if all evaluations pass, except for molecule availability.

**Selection.** The selection stage is the foundation of our evolutionary framework, ensuring the maintenance and progression of a population of candidate routes. In retrosynthesis planning, the success rate is commonly used to evaluate a route’s quality. However, within the evolutionary framework, most current routes are unsuccessful. Therefore, we introduce a partial reward mechanism based on SC score (Coley et al., 2018) to assess these incomplete routes. Given a route, we traverse it from the first step sequentially to identify the first invalid step. The molecule set at this step, denoted as  $\mathbb{M}$ , is then used to compute the reward for the route. The reward function  $F(\cdot)$  is defined as follows, where  $\mathbb{C}$  is the set of purchasable compounds:

$$F(\mathbb{M}) = - \sum_{m \in \mathbb{M}, m \notin \mathbb{C}} sc\_score(m)$$

The top  $n_c$  routes, as ranked by SC score are selected as the population for the next round of evolution.

**Mutation.** To optimize the current route, we explore the flexibility of LLMs in synthetic route reproduction. Specifically, we enable the LLM to analyze and edit the current route through prompt-based mutation. The LLM is instructed to modify the existing route or propose an alternative if deemed necessary, incorporating evaluation results from multiple perspectives as feedback. If the current route contains reaction-level errors, we retrieve reference routes from  $\mathbb{O}$ , weighted by their products’ Tanimoto similarity to the ‘product’ molecule in this step and provide them to the LLM.

Level	Type	Explanation
Molecule	Validity	Whether the molecule is valid (RDKit parsable)
	Availability	Whether the molecule is commercially available, i.e., in the building block stock
Reaction	Existence	Whether the reaction exists in the database
	Validity	Whether the product can be synthesized from the proposed reactants
Route	Connectivity	Whether the route is connected

Table 1: Three levels of feedback in the evaluation stage.

Algorithm	USPTO Easy			USPTO-190			Pistachio Reachable			Pistachio Hard		
	Solve Rate (%)			Solve Rate (%)			Solve Rate (%)			Solve Rate (%)		
	N=100	300	500	N=100	300	500	N=100	300	500	N=100	300	500
Graph2Edits(MCTS)	90.0	93.5	95.5	42.7	54.7	60.5	77.3	88.4	94.2	26.0	41.0	59.0
RootAligned(MCTS)	98.0	98.0	98.0	79.4	79.4	79.4	<b>99.3</b>	<b>99.3</b>	<b>99.3</b>	<b>83.0</b>	<b>83.0</b>	<b>83.0</b>
LocalRetro(MCTS)	92.5	94.5	95.5	44.3	50.9	58.3	86.7	90.0	95.3	52.0	55.0	62.0
Graph2Edits(Retro*)	92.0	95.5	97.0	51.1	59.4	78.5	94.0	95.0	95.5	71.0	74.0	79.0
RootAligned(Retro*)†	<b>99.0</b>	<b>99.0</b>	<b>99.0</b>	<b>86.8</b>	<b>86.8</b>	<b>86.8</b>	98.7	98.7	98.7	78.0	78.0	78.0
LocalRetro(Retro*)	95.5	97.5	98.0	51.0	65.8	73.7	97.3	<b>99.3</b>	<b>99.3</b>	63.0	69.0	72.0
LLM(MCTS)	54.5	68.5	75.5	25.8	27.2	31.3	12.7	17.3	20.7	0.0	4.0	5.0
LLM(Retro*)	56.0	69.0	75.5	23.2	26.8	30.6	14.7	19.3	13.3	0.0	2.0	5.0
LLM-Syn-Planner	<u>97.0</u>	<u>98.0</u>	<u>98.5</u>	<u>60.0</u>	<u>70.0</u>	<u>80.5</u>	92.0	94.7	96.7	64.0	73.0	80.0

Table 2: Summary of retrosynthesis planning performance across the four datasets. The best model for each experiment setting is bolded and the top three are underlined. All runs were limited to 30 minutes per molecule.  $N$  denotes the model call limit. † The RootAligned model does not achieve 100 model calls in 30 minutes due to high computational cost so the values across  $N$  are the same.

Additionally, for mutation queries, we retain the valid steps of the current route and provide the LLM with only the partial route starting from the first invalid step.

### 3.4 OPTIMIZATION FOR SYNTHESIZABLE MOLECULAR DESIGN

LLM-Syn-Planner can be easily extended to design optimized molecular structures alongside their corresponding synthesis pathways. A simple approach is to first optimize a molecular structure for the desired properties and then determine its synthesis pathway. As a proof of concept, we propose LLM-Syn-Designer, which integrates MolLEO (Wang et al., 2024) as the molecular structure optimizer, which leverages LLMs as genetic operators for molecular optimization. Specifically, we ask the LLM to generate a synthesizable molecule and analyze the synthetic route during the optimization process. To ensure synthesizability, we filter out molecules proposed by LLMs if their SC score exceeds 3.5 at each iteration of the optimization process. Additionally, in every round of evolutionary search, our framework acts as the synthesis pathway finder for the generated molecules. By combining these components, the integrated framework enables the end-to-end design of synthesizable molecules, harnessing the power of LLMs for both molecular optimization and synthesis planning.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETUP

**Dataset.** We conduct experiments using the USPTO (Schneider et al., 2016; Dai et al., 2019a) and Pistachio (pis) datasets. For USPTO, we utilize USPTO-190 (Chen et al., 2020b) and a simplified subset, USPTO-EASY, which is randomly sampled from the test set used in Retro\* single-step model training. For the Pistachio dataset, we adopt the version from Yu et al. (2024) but remove the starting material constraints. The route database is constructed using the training and validation sets from Retro\*, while the reaction database is a processed version of USPTO-Full, as used in Yu et al. (2024). For the building block set, we canonicalize all SMILES strings from the 23 million purchasable building blocks available in eMolecules, following the approach of Chen et al. (2020b). We show the statistics of the datasets in Appendix A.1.

**Baseline.** We consider three single-step retrosynthesis models in combination with two search algorithms: MCTS (Segler et al., 2017) and Retro\* (Chen et al., 2020b). The single-step models are: Graph2Edits (Zhong et al., 2023) is a template-free graph generative model that systematically edits the molecular graph of the target product to generate valid reactant structures. RootAligned (Zhong et al., 2022) is another template-free approach that enforces a strict one-to-one mapping between product and reactant SMILES strings by aligning them to a shared root atom. LocalRetro (Chen & Jung, 2021a) is a template-based method that employs local reaction templates involving atom and bond edits, coupled with a global attention mechanism to capture non-local effects.

**Metrics.** For retrosynthesis planning tasks, we use the success rate as the evaluation metric. For synthesizable molecular design tasks, we measure performance using the top-1 expected property in the designed molecules.






Figure 3: Fitness score of the best molecule found by each molecule optimization method. Only LLM-Syn-Designer here ensures the synthesizability of the found molecule.

**Configuration** We utilize GPT-4o<sup>1</sup> (Hurst et al., 2024) as our LLM and set the temperature to 0.7 for all queries, ensuring a balanced trade-off between creativity and reliability. To maintain efficiency, we impose a maximum search time of 30 minutes per molecule.

#### 4.2 RETROSYNTHESIS PLANNING

We present the retrosynthesis planning results in Table 2. The LLM-based approaches show a clear distinction between using LLMs as single-step predictors within a search algorithm and leveraging them to generate complete retrosynthetic routes optimized via tree evolutionary algorithms (LLM-Syn-Planner). When LLMs are integrated into MCTS or Retro\*, their solve rates are significantly lower than those of traditional models, particularly on challenging datasets (e.g., Pistachio Hard, where solve rates are near zero). This suggests that current LLM-based single-step models struggle to produce high-quality reaction predictions, leading to suboptimal search performance. Moreover, increasing the number of model calls does not consistently improve results, especially on the USPTO-190 and Pistachio datasets, highlighting intrinsic limitations in LLMs’ single-step reaction prediction capabilities.

In contrast, LLM-Syn-Planner performs remarkably well, achieving solve rates comparable to—or even exceeding—some single-step model-guided search. Notably, LLM-Syn-Planner significantly outperforms LLM (MCTS/Retro\*), indicating that optimizing full multi-step retrosynthetic routes rather than predicting step-by-step transformations enhances LLM effectiveness. While LLMs may not yet rival expert-designed single-step models in reaction prediction precision, they can generate promising retrosynthetic routes by using their long-term planning capabilities.

These findings suggest that while LLMs are not yet competitive as single-step predictors in MCTS/Retro\*, their strength lies in generating full retrosynthetic pathways that can be optimized through evolutionary techniques. This underscores a potential shift in focus from improving LLMs for single-step retrosynthesis to developing methods that exploit their generative capabilities for full-route planning combined with downstream optimization strategies like EA. We show ablation studies in Appendix C.1.

#### 4.3 SYNTHESIZABLE DESIGN

To evaluate the synthesizable design capability of LLM, we first consider common heuristic oracle functions relevant to bioactivity and drug discovery. We compare LLM-Syn-Designer with various molecular optimization methods, including Graph-GA Jensen (2019), REINVENT Olivecrona et al. (2017), MolLEO Wang et al. (2024), and MARS Xie et al. (2021), and present the results in Figure 3. Notably, the baseline methods do not enforce synthesizability constraints, allowing them to explore a broader chemical space and achieve higher scores, albeit with non-synthesizable molecules. The results demonstrate that LLM-Syn-Designer effectively balances optimization efficiency and synthesizability. In all cases, the best molecules identified by LLM-Syn-Designer exhibit competitive or superior fitness compared to traditional algorithms and MolLEO, while ensuring synthesizability. Specifically, for the isomers\_C9H10N2O2PF2Cl target, LLM-designs

<sup>1</sup>GPT-4o-2024-11-20



achieve comparable or higher scaled fitness values with fewer oracle calls than all other methods. This suggests that integrating synthesizability constraints within the optimization process does not necessarily compromise efficiency.

## 5 RELATED WORK

### 5.1 MACHINE LEARNING SINGLE-STEP RETROSYNTHESIS MODEL

Single-step retrosynthesis models predict the outcome of a single reaction step, i.e., given an input molecule, how can it be decomposed and into which constituents? Early works directly predicted precursors by seq-to-seq translation on SMILES Weininger (1988); Liu et al. (2017) or using fingerprints Segler & Waller (2017); Coley et al. (2017); Fortunato et al. (2020). More recently, single-step retrosynthesis models have employed transformers Vaswani (2017) and graph neural networks (GNNs). Methods can be categorized into template-based, template-free, or semi-template methods. Template-based methods use pre-defined chemical rules which can be advantageous if they are defined with high granularity Szymkuć et al. (2016); Grzybowski et al. (2018); Segler & Waller (2017); Dai et al. (2019b); Ishida et al. (2019); Seidl et al. (2022); Chen & Jung (2021b); Xie et al. (2023). Template-free methods attempt to learn these rules from data and learn a translation Liu et al. (2017); Zheng et al. (2019); Schwaller et al. (2020); Zhong et al. (2022). Finally, semi-template methods make intermediate predictions (such as synthons) and then predict the precursors based on these Shi et al. (2020); Somnath et al. (2021); Sacha et al. (2021); Zhong et al. (2023).

### 5.2 SEARCH-DIRECTED RETROSYNTHESIS PLANNING

By coupling a search algorithm with single-step retrosynthesis models, multi-step retrosynthesis can be performed. Exemplary works include applying Monte Carlo tree search (MCTS) Segler et al. (2018a), Retro\* Chen et al. (2020b), Planning with Dual Value Networks (PDVN) Liu et al. (2023), and a recent double-ended search algorithm Yu et al. (2024). Since retrosynthesis has broad applicability for molecular discovery, many retrosynthesis platforms exist, encompassing industrial Szymkuć et al. (2016); Grzybowski et al. (2018); Genheden et al. (2020); Saigiridharan et al. (2024); Watson et al. (2019); Molecule.one; Schwaller et al. (2020) and open-source Genheden et al. (2020); Saigiridharan et al. (2024); Coley et al. (2019); Tu et al. (2025) solutions. Very recently, works have investigated applying LLMs for retrosynthesis through fine-tuning Nguyen-Van et al. (2024), instruction-tuning Yang et al. (2024), platform assistants Zhang et al. (2025), experimental planning agents Liu et al. (2024b), and integration with knowledge graphs for synthesis planning of polymers Ma et al. (2025).

## 6 CONCLUSION

In this paper, we studied the retrosynthesis problem with LLMs. Specifically, we experimented with using LLMs as single-step reaction prediction models with a search algorithm and found LLMs significantly underperformed specialized reaction models. To improve this, we proposed to sample entire multi-step synthetic pathways and introduced an evolutionary process to optimize them. To scale this approach, we leveraged a linear format to store reaction steps and designed partial rewards with retrieved reaction sub-trajectories. In the end, we bridged the performance gap and matched the SOTA performance in retrosynthesis planning. In addition, we demonstrated LLMs can be easily adapted to the synthesizable molecular design problem to find property-optimized molecules that are synthesizable.

**Limitation and future work:** Despite promising results, we observed that LLMs suffered significantly with sparse rewards (e.g. in the shooting setup) while improved significantly with partial rewards and retrieved sub-trajectories. It is worth studying how to incorporate a search algorithm into our framework when LLMs struggle to generate any synthesis paths with the desired target molecule. For future work, it is promising to study more flexible design criteria enabled by LLMs such as material-constrained synthesis planning (Guo & Schwaller, 2024b).

## ETHICS STATEMENT

This work investigates how LLMs can be used for retrosynthesis planning and synthesizable molecule design. Both use-cases are applicable to therapeutics and materials design. No molecules were synthesized and experimentally tested so there are no specific societal consequences we feel should be highlighted. However, in the future, the framework, if properly experimentally validated, could have positive societal benefits.

## REFERENCES

- Pistachio, january 2024. URL <https://www.nextmovesoftware.com/pistachio.html>.
- David C Blakemore, Luis Castro, Ian Churcher, David C Rees, Andrew W Thomas, David M Wilson, and Anthony Wood. Organic synthesis provides opportunities to transform drug discovery. *Nature chemistry*, 10(4):383–394, 2018.
- Daniil A Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- Jonas Boström, Dean G Brown, Robert J Young, and György M Keserü. Expanding the medicinal chemistry synthetic toolbox. *Nature Reviews Drug Discovery*, 17(10):709–727, 2018.
- John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato. A model to search for synthesizable molecules. *Advances in Neural Information Processing Systems*, 32, 2019.
- John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato. Barking up the right tree: an approach to search over molecule synthesis dags. *Advances in neural information processing systems*, 33:6852–6866, 2020.
- Liao Chang, Wenhao Zheng, Huaxiu Yao, and Ying Wei. How well can llms synthesize molecules? an llm-powered framework for multi-step retrosynthesis.
- Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro\*: learning retrosynthetic planning with neural guided a\* search. In *International conference on machine learning*, pp. 1608–1616. PMLR, 2020a.
- Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro\*: learning retrosynthetic planning with neural guided a\* search. In *International conference on machine learning*, pp. 1608–1616. PMLR, 2020b.
- Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021a.
- Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021b.
- 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.
- 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.
- Connor W Coley, Dale A Thomas III, Justin AM Lummiss, Jonathan N Jaworski, Christopher P Breen, Victor Schultz, Travis Hart, Joshua S Fishman, Luke Rogers, Hanyu Gao, et al. A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365(6453):eaax1566, 2019.
- 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.

- Elias James Corey, Alan K Long, and Stewart D Rubenstein. Computer-assisted analysis in organic synthesis. *Science*, 228(4698):408–418, 1985.
- Miruna Cretu, Charles Harris, Julien Roy, Emmanuel Bengio, and Pietro Liò. Synflownet: Towards molecule design with guaranteed synthesis pathways. *arXiv preprint arXiv:2405.01155*, 2024.
- 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, 2019a.
- 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, 2019b.
- Inc. Daylight Chemical Information Systems. 4. smarts - a language for describing molecular patterns.
- Michael E Fortunato, Connor W Coley, Brian C Barnes, and Klavs F Jensen. Data augmentation and pretraining for template-based retrosynthetic prediction in computer-aided synthesis planning. *Journal of chemical information and modeling*, 60(7):3398–3407, 2020.
- Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. *Proc. 10th International Conference on Learning Representations*, 2022.
- Wenhao Gao, Shitong Luo, and Connor W Coley. Generative artificial intelligence for navigating synthesizable chemical space. *arXiv preprint arXiv:2410.03494*, 2024.
- 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.
- Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao Liu, Simon Blackburn, Karam Thomas, Connor Coley, Jian Tang, et al. Learning to navigate the synthetically accessible chemical space using reinforcement learning. In *International conference on machine learning*, pp. 3668–3679. PMLR, 2020.
- Bartosz A Grzybowski, Sara Szymkuć, Ewa P Gajewska, Karol Molga, Piotr Dittwald, Agnieszka Wołos, and Tomasz Klucznik. Chematica: a story of computer code that started to think like a chemist. *Chem*, 4(3):390–398, 2018.
- Jeff Guo and Philippe Schwaller. Directly optimizing for synthesizability in generative molecular design using retrosynthesis models. *arXiv preprint arXiv:2407.12186*, 2024a.
- Jeff Guo and Philippe Schwaller. It takes two to tango: Directly optimizing for constrained synthesizability in generative molecular design. *arXiv preprint arXiv:2410.11527*, 2024b.
- Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical space via deep reinforcement learning. *ACS omega*, 5(51):32984–32994, 2020.
- Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. Understanding the planning of llm agents: A survey. *arXiv preprint arXiv:2402.02716*, 2024.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Ilia Igashov, Arne Schneuing, Marwin Segler, Michael M Bronstein, and Bruno Correia. Retrobridge: Modeling retrosynthesis with markov bridges. In *The Twelfth International Conference on Learning Representations*.
- Shoichi Ishida, Kei Terayama, Ryosuke Kojima, Kiyosei Takasu, and Yasushi Okuno. Prediction and interpretable visualization of retrosynthetic reactions using graph convolutional networks. *Journal of chemical information and modeling*, 59(12):5026–5033, 2019.
- Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.

- Zygimantas Jocys, Henriette MG Willems, and Katayoun Farrahi. Synthformer: Equivariant pharmacophore-based generation of molecules for ligand-based drug design. *arXiv preprint arXiv:2410.02718*, 2024.
- Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, pp. 3393–3403. PMLR, 2020.
- Michał Koziarski, Andrei Rekesh, Dmytro Shevchuk, Almer van der Sloot, Piotr Gaiński, Yoshua Bengio, Cheng-Hao Liu, Mike Tyers, and Robert A Batey. Rgfn: Synthesizable molecular generation using gflownets. *arXiv preprint arXiv:2406.08506*, 2024.
- Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS central science*, 3(10):1103–1113, 2017.
- Gang Liu, Michael Sun, Wojciech Matusik, Meng Jiang, and Jie Chen. Multimodal large language models for inverse molecular design with retrosynthetic planning. *arXiv preprint arXiv:2410.04223*, 2024a.
- 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*, pp. 22266–22276. PMLR, 2023.
- Xuan Liu, Solana Chiu, and Huimin Zhao. Autonomous faithful retrosynthesis with large language models: From synthesis planning to experimental procedures. In *2024 AIChE Annual Meeting*. AIChE, 2024b.
- Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W Coley, and Jianzhu Ma. Projecting molecules into synthesizable chemical spaces. *Proc. 41st International Conference on Machine Learning*, 2024.
- Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, pp. 1–11, 2024.
- Qinyu Ma, Yuhao Zhou, and Jianfeng Li. Leveraging large language models as knowledge-driven agents for reliable retrosynthesis planning. *arXiv preprint arXiv:2501.08897*, 2025.
- Krzysztof Maziarz, Austin Tripp, Guoqing Liu, Megan Stanley, Shufang Xie, Piotr Gaiński, Philipp Seidl, and Marwin Segler. Re-evaluating retrosynthesis algorithms with syntheseus. *arXiv preprint arXiv:2310.19796*, 2023.
- Molecule.one. The m1 platform.
- Phuong Nguyen-Van, Long Nguyen Thanh, Ha Hoang Manh, Ha Anh Pham Thi, Thanh Le Nguyen, and Viet Anh Nguyen. Adapting language models for retrosynthesis prediction. 2024.
- Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9:1–14, 2017.
- 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.
- Lakshidaa Saigiridharan, Alan Kai Hassen, Helen Lai, Paula Torren-Peraire, Ola Engkvist, and Samuel Genheden. Aizynthfinder 4.0: developments based on learnings from 3 years of industrial application. *Journal of Cheminformatics*, 16(1):57, 2024.
- Nadine Schneider, Nikolaus Stiefl, and Gregory A Landrum. What’s what: The (nearly) definitive guide to reaction role assignment. *Journal of chemical information and modeling*, 56(12):2336–2346, 2016.

- Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical science*, 11(12):3316–3325, 2020.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*, 2024.
- 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.
- Marwin HS Segler, Mike Preuss, and Mark P Waller. Learning to plan chemical syntheses. *arXiv preprint arXiv:1708.04202*, 2017.
- Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018a.
- Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018b.
- Philipp Seidl, Philipp Renz, Natalia Dyubankova, Paulo Neves, Jonas Verhoeven, Jorg K Wegner, Marwin Segler, Sepp Hochreiter, and Gunter Klambauer. Improving few- and zero-shot reaction template prediction using modern hopfield networks. *Journal of chemical information and modeling*, 62(9):2111–2120, 2022.
- Seonghwan Seo, Minsu Kim, Tony Shen, Martin Ester, Jinkyoo Park, Sungsoo Ahn, and Woo Youn Kim. Generative flows on synthetic pathway for drug design. *arXiv preprint arXiv:2410.04542*, 2024.
- 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*, pp. 8818–8827. PMLR, 2020.
- 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.
- Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2998–3009, 2023.
- Kyle Swanson, Gary Liu, Denise B Catacutan, Autumn Arnold, James Zou, and Jonathan M Stokes. Generative ai for designing and validating easily synthesizable and structurally novel antibiotics. *Nature Machine Intelligence*, 6(3):338–353, 2024.
- Sara Szymkuć, Ewa P Gajewska, Tomasz Klucznik, Karol Molga, Piotr Dittwald, Michał Startek, Michał Bajczyk, and Bartosz A Grzybowski. Computer-assisted synthetic planning: the end of the beginning. *Angewandte Chemie International Edition*, 55(20):5904–5937, 2016.
- Zhengkai Tu, Sourabh J Choure, Mun Hong Fong, Jihye Roh, Itai Levin, Kevin Yu, Joonyoung F Joung, Nathan Morgan, Shih-Cheng Li, Xiaoqi Sun, et al. Askcos: an open source software suite for synthesis planning. *arXiv preprint arXiv:2501.01835*, 2025.
- A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Haorui Wang, Marta Skreta, Cher-Tian Ser, Wenhao Gao, Lingkai Kong, Felix Strieth-Kalthoff, Chenru Duan, Yuchen Zhuang, Yue Yu, Yanqiao Zhu, et al. Efficient evolutionary search over chemical space with large language models. *arXiv preprint arXiv:2406.16976*, 2024.
- Ian A Watson, Jibo Wang, and Christos A Nicolaou. A retrosynthetic analysis algorithm implementation. *Journal of cheminformatics*, 11:1–12, 2019.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- 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.
- 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, pp. 5330–5338, 2023.
- Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. Mars: Markov molecular sampling for multi-objective drug discovery. In *Proc. 9th International Conference on Learning Representations*, 2021.
- Yifei Yang, Runhan Shi, Zuchao Li, Shu Jiang, Bao-Liang Lu, Yang Yang, and Hai Zhao. Batgpt-chem: A foundation large model for retrosynthesis prediction. *arXiv preprint arXiv:2408.10285*, 2024.
- Kevin Yu, Jihye Roh, Ziang Li, Wenhao Gao, Runzhong Wang, and Connor W Coley. Double-ended synthesis planning with goal-constrained bidirectional search. *arXiv preprint arXiv:2407.06334*, 2024.
- Chonghuan Zhang, Qianghua Lin, Biwei Zhu, Haopeng Yang, Xiao Lian, Hao Deng, Jiajun Zheng, and Kuangbiao Liao. Synask: unleashing the power of large language models in organic synthesis. *Chemical Science*, 16(1):43–56, 2025.
- Zirui Zhao, Wee Sun Lee, and David Hsu. Large language models as commonsense knowledge for large-scale task planning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Shuangjia Zheng, Jiahua Rao, Zhongyue Zhang, Jun Xu, and Yuedong Yang. Predicting retrosynthetic reactions using self-corrected transformer neural networks. *Journal of chemical information and modeling*, 60(1):47–55, 2019.
- 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.
- Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Min Wu, Tingjun Hou, and Mingli Song. Root-aligned smiles: a tight representation for chemical reaction prediction. *Chemical Science*, 13(31):9023–9034, 2022.
- Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Tingjun Hou, and Mingli Song. Recent advances in deep learning for retrosynthesis. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 14(1):e1694, 2024.
- Yuchen Zhuang, Xiang Chen, Tong Yu, Saayan Mitra, Victor Bursztyn, Ryan A Rossi, Somdeb Sarkhel, and Chao Zhang. Toolchain\*: Efficient action space navigation in large language models with a\* search. *arXiv preprint arXiv:2310.13227*, 2023.

Name	No. of Routes	Avg. Route Length	Avg. SA score	Avg. SC score
USPTO Easy	200	3.7	2.8	3.8
USPTO-190	190	6.7	3.6	4.0
Pistachio Reachable	150	5.5	3.1	3.9
Pistachio Hard	100	7.5	3.6	3.9

Table 3: Statistics of the dataset used in the experiments.

## A EXTENDED DESCRIPTIONS

### A.1 DATASET STATISTICS

We show the dataset statistics in Table 3.

### A.2 MCTS FOR RETROSYNTHESIS PLANNING

The single-step model predicts potential sets of reactants for a given product, transforming a target molecule into plausible precursors. However, multiple steps may be needed to reach commercially available or easily synthesized materials. This is why the single-step reaction model is integrated with MCTS: it systematically explores these multi-step routes, pruning unlikely paths while focusing on the most promising transformations. By striking a balance between exploration and exploitation, MCTS avoids getting stuck in unproductive branches and can uncover synthetic routes that might not be obvious through manual inspection alone.

Under the MCTS procedure, the target molecule is defined as the root node of a search tree, and each edge represents a single-step retrosynthetic transformation predicted by the reaction model. A policy network can be used to rank or filter the most promising disconnection suggestions at each step, while a value function provides an estimate of how likely a given partial route is to succeed in the long run. The algorithm selects which node to expand next using an Upper Confidence Bound (UCB), which balances the value estimate (exploitation) with the uncertainty in that estimate (exploration). A reward function then quantifies the outcome of each expansion—often based on reaction feasibility, synthetic cost, or reaching known starting materials. These reward signals are backpropagated to update the value estimates of each node. Finally, iterating selection, expansion, simulation, and backpropagation until we reach a termination condition (time limit, enough solutions found).

### A.3 RETRO\* ALGORITHM

Retro\* (Chen et al., 2020b) integrates neural networks with a best-first search strategy to solve retrosynthesis problems. It models the problem as an AND-OR tree, where "AND" nodes represent reactions and "OR" nodes correspond to molecules. A neural network, trained on prior retrosynthesis experiences, estimates the cost of each node. Using a best-first search, the algorithm prioritizes the most promising pathways based on these predictions. It then applies a single-step model to expand the selected node, generating an AND-OR subtree. Finally, it updates the pathway costs to guide the next selection step.

### A.4 ADDITIONAL EXPERIMENTAL DETAILS

For single-step models, we use the checkpoints from synthesus<sup>2</sup>. In the MCTS algorithm, we employ a basic reward function: a state receives a reward of 1.0 if all molecules are purchasable (i.e., the state is solved), and 0.0 otherwise. The value function is set as a constant 0.5. For policy, we use softmax values derived from the single-step reaction model, scaled by a temperature of 3.0 and normalized across the total number of reactions.

In the Retro\* algorithm, we follow the retro\*-0 variant described in the original paper (Chen et al., 2020b). The OrNode cost function assigns a cost of 0 to purchasable molecules and infinity otherwise. The AndNode cost function defines the reaction cost as  $-\log(\text{softmax})$  of the reaction model output, thresholded at a minimum value. For the search heuristic (value function), we use a constant value of 0, consistent with the retro\*-0 algorithm.

<sup>2</sup><https://github.com/microsoft/synthesus>



Format	USPTO Easy	USPTO 190
Textual + Extraction	71.5	27.9
Tree	54.5	12.1
Sequential	97.0	60.0

Table 4: Ablation study of using different route format in syn-planner.

Format	USPTO Easy	USPTO 190
w/o RAG	55.0	16.3
w/ RAG	97.0	60.0

Table 5: Comparison of using and without using molecule RAG in the INITIALIZATION and MUTATION prompt.

## A.5 COMPUTATIONAL RESOURCES

Our experiments utilized the GPT-4o model; this refers to the GPT-4o checkpoint from 2024-11-20<sup>3</sup>. All GPT-4o checkpoints were hosted on Microsoft Azure<sup>4</sup>.

## B EXTENDED RELATED WORK

### B.1 SYNTHESIZABLE MOLECULAR DESIGN

Synthesizable molecular design aims to generate molecules (with optimal properties) that are also synthesizable, as predicted by a retrosynthesis model. While retrosynthesis methods are often described as "top-down" because they decompose a target molecule into purchasable precursors, the most common methods in literature for synthesizable molecular design proceeds "bottom-up", which combine building blocks to construct the final molecule. Therefore, instead of predicting the resulting *precursors* from an input molecule, "bottom-up" approaches require a way to predict the *product molecule* given precursors. To this end, existing approaches either use forward synthesis prediction models Bradshaw et al. (2019; 2020) or define a set of templates which dictate *how* building blocks can be combined Gao et al. (2022; 2024); Luo et al. (2024); Koziarski et al. (2024); Cretu et al. (2024); Seo et al. (2024); Swanson et al. (2024); Jocys et al. (2024). These methods can be broadly classified as synthesizability-constrained generative models. An alternative approach is to couple retrosynthesis models directly into the optimization loop of generative models, such that synthesizability is optimized for, rather than enforced in the generation process Guo & Schwaller (2024a;b).

## C EXTENDED EXPERIMENT RESULTS

### C.1 ABLATION STUDY

#### Observation 1: The linear format of synthesis steps significantly outperforms the tree format.

We investigate the influence of route format in Table 4. The results suggest that linear storage of decision steps better reduces the exponentially growing complexity of the synthesis pathway, thus leading to much higher success rates. Additionally, we introduce a simple baseline named (Textual + Extraction) to allow the LLM to generate in an arbitrary format, followed by a subsequent query to extract the route from the returned response. Surprisingly, this approach also yields decent performance, even with an unconstrained format.

**Observation 2: Even rough intermediate feedback can be significantly useful for LLMs.** To assess the impact of incorporating Molecule RAG into the retrosynthesis planning process, we remove RAG in the `INITIALIZATION` and `MUTATION` prompts. As summarized in Table 5, even if the Morgan fingerprint does not directly reflect synthesis similarity and similar molecular structures do not always exist in the database, RAG still significantly improves the performance.

<sup>3</sup><https://platform.openai.com/docs/models>

<sup>4</sup>[\\*.openai.azure.com](https://openai.azure.com)

	USPTO Easy	USPTO 190
w/ partial reward	97.0	60.0
w/ only final reward	68.0	22.6

Table 6: Comparison of using and without using partial rewards in the selection stage of syn-planner.

**Observation 3: Partial reward is crucial for long-horizon sequential decision-making.** The target reward is very sparse as it only evaluates if a generated synthesis pathway is valid. We validate the importance of the partial reward by a simple synthesis accessibility evaluator (SC score) by comparing it against one with only the final reward. With partial reward, the success rate improves considerably across both datasets.

## C.2 CASE STUDY




Figure 4: Top 1 molecule of gsk3 $\beta$  found by LLM-Syn-Designer.




Figure 5: Top 1 molecule of jnk3 found by LLM-Syn-Designer.

## D PROMPTS

We show the prompts of **INITIALIZATION** and **MUTATION** for LLM-Syn-Planner. And LLM operators prompt for LLM-Syn-Designer.

## LLM-Syn-Planner INITIALIZATION prompts

As a professional chemist specialized in synthesis analysis, you are tasked with generating a retrosynthesis route for a target molecule provided in SMILES format.

A retrosynthesis route is a series of retrosynthesis steps that starts from the target molecule and ends with some commercially purchasable compounds. The reactions are from the USPTO dataset. Please also consider reactions in stereochemistry.

The route should be a list of steps wrapped in <ROUTE></ROUTE> with <EXPLANATION></EXPLANATION> after it. Each step in the list should be a dictionary. You need to keep a molecule set, which consists of the molecules we need to synthesize or purchase. In each step, you need to select a molecule from the 'Molecule set' as the product molecule in this step and use a backward reaction to find the reactants. After taking the backward reaction in this step, you need to remove the product molecule from the molecule set and add the reactants you find into the molecule set, and then name this updated set as the 'Updated molecule set' in this step. In the next step, the starting molecule set should be the 'Updated molecule set' from the previous step. In the last step, all the molecules in the 'Updated molecule set' should be purchasable. Here is an example:

```
<ROUTE>
[
  {
    'Molecule set': "[Target Molecule]",
    'Rational': Step analysis,
    'Product': "[Product molecule]",
    'Reaction': "[Reaction template]",
    'Reactants': "[Reactant1, Reactant2]",
    'Updated molecule set': "[Reactant1, Reactant2]"
  },
  {
    'Molecule set': "[Reactant1, Reactant2]",
    'Rational': Step analysis,
    'Product': "[Product molecule]",
    'Reaction': "[Reaction template]",
    'Reactants': "[subReactant1, subReactant2]",
    'Updated molecule set': "[Reactant1, subReactant1, subReactant2]"
  }
]
</ROUTE>
<EXPLANATION>: Explanation for the whole route. </EXPLANATION>\\
```

## Requirements:

1. The 'Molecule set' contains molecules we need to synthesize at this stage. In the first step, it should be the target molecule. In the following steps, it should be the 'Updated molecule set' from the previous step.
  2. The 'Rational' part in each step should be your analysis for synthesis planning in this step. It should be in the string format wrapped with ''
  3. 'Product' is the molecule we plan to synthesize in this step. It should be from the 'Molecule set'. The molecule should be a molecule from the 'Molecule set' in a list. The molecule smiles should be wrapped with ''
  4. 'Reaction' is a reaction that can synthesize the product molecule. It should be on a list. The reaction template should be in SMILES format. For example, [Product]>>Reactant1.Reactant2].
  5. 'Reactants' are the reactants of the reaction. It should be on a list. The molecule smiles should be wrapped with ''
  6. The 'Updated molecule set' should be molecules we need to purchase or synthesize after taking this reaction. To get the 'Updated molecule set', you need to remove the product molecule from the 'Molecule set' and then add the reactants in this step into it. In the last step, all the molecules in the 'Updated molecule set' should be purchasable.
  7. In the <EXPLANATION>, you should analyze the whole route and ensure the molecules in the 'Updated molecule set' in the last step are all purchasable.
- My target molecule is:  
{Target Molecule}
- To assist you, example retrosynthesis routes that are either close to the target molecule or representative will be provided.

```
<ROUTE>
Retrieved route here
</ROUTE>
```

Please propose a retrosynthesis route for my target molecule. The provided reference routes may be helpful. You can also design a synthetic route based on your own knowledge.

## LLM-Syn-Planner MUTATION prompts

As a professional chemist specializing in synthesis analysis, you are tasked with modifying a retrosynthesis route for target molecules provided in SMILES format. A retrosynthesis route is a series of retrosynthesis steps that starts from the given target molecule set and ends with some commercially purchasable compounds. In the route, you need to keep a molecule set, which are the molecules we need. In the first step, the molecule set should be the target molecule set given by the user. In each step, you need to provide a backward reaction and update this molecule set. Specifically, you need to remove the product molecule of the reaction from the molecule set and then add the reactants to it.

By doing so, you will end with a molecule set in which all the molecules are commercially purchasable. The reactions are from the USPTO dataset. Please also take reactions in stereochemistry into consideration. For example, E-configuration or Z-configuration.

The route should be a list of steps wrapped in <ROUTE></ROUTE> with <EXPLANATION></EXPLANATION> after it. Each step in the list should be a dictionary. You need to keep a molecule set, which consists of the molecules we need to synthesize or purchase. In each step, you need to select a molecule from the 'Molecule set' as the product molecule in this step and use a backward reaction to find the reactants. After taking the backward reaction in this step, you need to remove the product molecule from the molecule set and add the reactants you find into the molecule set, and then name this updated set as the 'Updated molecule set' in this step. In the next step, the starting molecule set should be the 'Updated molecule set' from the previous step. In the last step, all the molecules in the 'Updated molecule set' should be purchasable. Here is an example:

```
<ROUTE>
[
  {
    'Molecule set': "[Target Molecule]",
    'Rational': 'Step analysis',
    'Product': "[Product molecule]",
    'Reaction': "[Reaction template]",
    'Reactants': "[Reactant1, Reactant2]",
    'Updated molecule set': "[Reactant1, Reactant2]"
  },
  {
    'Molecule set': "[Reactant1, Reactant2]",
    'Rational': 'Step analysis',
    'Product': "[Product molecule]",
    'Reaction': "[Reaction template]",
    'Reactants': "[subReactant1, subReactant2]",
    'Updated molecule set': "[Reactant1, subReactant1, subReactant2]"
  }
]
</ROUTE>
```

<EXPLANATION>: Explanation for the whole route. </EXPLANATION>\\

Requirements:

1. The 'Molecule set' contains molecules we need to synthesize at this stage. In the first step, it should be the target molecule set. In the following steps, it should be the 'Updated molecule set' from the previous step.
2. The 'Rational' part in each step should be your analysis for synthesis planning in this step. It should be in the string format wrapped with ''.
3. 'Product' is the molecule we plan to synthesize in this step. It should be from the 'Molecule set'. The molecule should be a molecule from the 'Molecule set' in a list. The molecule smiles should be wrapped with ''.
4. 'Reaction' is a reaction that can synthesize the product molecule. It should be on a list. The reaction template should be in SMILES format. For example, [Product>Reactant1.Reactant2].
5. 'Reactants' are the reactants of the reaction. It should be on a list. The molecule smiles should be wrapped with ''.
6. The 'Updated molecule set' should be molecules we need to purchase or synthesize after taking this reaction. To get the 'Updated molecule set', you need to remove the product molecule from the 'Molecule set' and then add the reactants in this step into it. In the last step, all the molecules in the 'Updated molecule set' should be purchasable.
7. In the <EXPLANATION>, you should analyze the whole route and ensure the molecules in the 'Updated molecule set' in the last step are all purchasable.

My target molecule set is:

```
{Target Molecule set}
```

Here are the feedbacks for the route:

```
{Feedback}
```

To assist you, example retrosynthesis routes that are close to the target molecules in the starting molecule set will be provided.

```
<ROUTE>
```

Retrieved route here  
</ROUTE>

Please propose a retrosynthesis route for the starting molecule set. The provided reference routes may be helpful. You can also design a synthetic route based on your own knowledge. All the molecules should be in SMILES format. For example, Cl2 should be ClCl in SMILES format. Br2 should be BrBr in SMILES format. H2O should be O in SMILES format. HBr should be [H]Br in SMILES format. NH3 should be N in SMILES format. Hydrogen atoms are implicitly understood unless explicitly needed for clarity.

#### LLM-Syn-Designer prompts

I have two molecules and their JNK3 scores. The JNK3 score measures a molecule's biological activity against JNK3.  
Molecule 1 SMILES, Molecule 1 score  
Molecule 2 SMILES, Molecule 2 score

Now I want to synthesize a new molecule that has a higher JNK3 score. Please propose a new synthesizable molecule that has a higher JNK3 score. You can either make crossover and mutations based on the given molecules or just propose a new molecule based on your knowledge.

Your output should follow the format:

<EXPLANATION>Your analysis</EXPLANATION>  
<MOLECULE>The SMILES of your proposed molecule</MOLECULE>

Here are the requirements:

1. In the <EXPLANATION>, you should analyze how to edit the given molecules to get a better property score and then propose your edited molecule or your proposed new molecule, and how to synthesize your proposed/edited molecule.
2. In the <MOLECULE>, you should provide the SMILES of the molecule you propose.