

Digital Discovery

Accepted Manuscript

This article can be cited before page numbers have been issued, to do this please use: S. Prakash, N. Patel, H. Jacobsen and V. K. Prasad, *Digital Discovery*, 2026, DOI: 10.1039/D5DD00562K.



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

A USER-TUNABLE MACHINE LEARNING FRAMEWORK FOR STEP-WISE SYNTHESIS PLANNING

Shivesh Prakash¹  Nandan Patel⁴  Hans-Arno Jacobsen^{1,2}  Viki Kumar Prasad^{2,3,4*} 

¹Department of Computer Science, University of Toronto, 40 St George St, Toronto, ON M5S 2E4

²The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto, 10 King's College Rd, Toronto, ON M5S 3G8

³Data Sciences Institute, University of Toronto, 700 University Ave 10th floor, Toronto, ON M7A 2S4

⁴Current affiliation: Department of Chemistry, University of Calgary, 2500 University Drive NW, Calgary, AB T2N 1N4

*vikikumar.prasad@ucalgary.ca

ABSTRACT

We introduce MHNpath, a machine learning-driven retrosynthetic tool designed for computer-aided synthesis planning. Leveraging modern Hopfield networks and novel comparative metrics, MHNpath efficiently prioritizes reaction templates, improving the scalability and accuracy of retrosynthetic predictions. The tool incorporates a tunable scoring system that allows users to prioritize pathways based on cost, reaction temperature, and toxicity, thereby facilitating the design of greener and cost-effective reaction routes. We demonstrate its effectiveness through case studies involving complex molecules from ChemByDesign, showcasing its ability to predict novel synthetic and enzymatic pathways. Furthermore, we benchmark MHNpath against existing frameworks using the PaRoutes dataset, achieving a solution rate of 85.4% and replicating 69.2% of experimentally validated “gold-standard” pathways. Our case studies show that the tool can generate shorter, cheaper, moderate-temperature routes employing green solvents, as exemplified by molecules such as dronabinol, arformoterol, and lupinine.

1 Introduction

The integration of machine learning (ML) into chemical synthesis has transformed the field of computer-aided synthesis planning (CASP), providing chemists with powerful tools to design and execute synthetic routes more efficiently. Traditional retrosynthetic analysis, which relies heavily on expert intuition and experience^{1,2}, is increasingly being augmented by data-driven approaches³ to break down complex molecules into simpler precursors using ML techniques to predict viable synthetic pathways^{1,4}. These advancements are particularly crucial in addressing the growing complexity of modern synthetic challenges, where exploring vast chemical spaces catalyzes the discovery of novel molecules. Examples of such challenges include the synthesis of complex natural products with multiple stereocenters⁵ and the need for highly selective functional group transformations in pharmaceutical development⁶.

In CASP, clean chemical data is the foundation for computationally designing synthetic routes. These data enables training of a predictive model that finds precursors while an efficient search algorithm navigates the expansive chemical space to propose feasible retrosynthetic pathways. State-of-the-art ML models for CASP can be categorized into template-based and template-free approaches where templates are generalized representations that encapsulate the core chemical transformation patterns inherent to each reaction. Complementing the ML model and search algorithm is an integrated scoring system that evaluates routes based on multiple criteria, such as cost and reaction temperature. Together, these three components establish a cohesive framework that underpins modern computer-aided synthesis planning.

Recent advancements in ML, including the development of deep neural networks and transformer-based models, have significantly enhanced the predictive capabilities of template-free methods in CASP^{7–11}. Transformer models have emerged as a powerful tool in retrosynthetic prediction due to their ability to handle sequence-to-sequence tasks



effectively¹². The Molecular Transformer model⁸, and its variants^{9–11}, for instance, has demonstrated remarkable accuracy in predicting both reactants and reaction conditions by treating retrosynthesis as a language translation problem from products to reactants. These models leverage self-attention mechanisms to capture complex dependencies within chemical reactions, making them highly effective for single-step retrosynthesis^{13,14}. Some graph-based methods^{15,16} predict the products from reactants in an auto-regressive fashion, while others^{17–19} predict the products by anticipating the final states of bonds or electrons. However, the reliance of template-free approaches on extensive labeled reaction data, computational resources, and single-step focus has been a significant limitation in CASP.

On the other hand, template-based approaches have also been fundamental in CASP.²⁰ These approaches directly apply predefined reaction rules to fragment target molecules into simpler structures, making them more interpretable and chemically intuitive compared to template-free approaches. Segler et al.²¹ and Coley et al.²² were among the first to use automatically extracted rules to predict outcomes of organic reactions using neural networks while Zhang et al.²³ extend it by also continually evolving them and Rho et al.²⁴ give a global outlook by effectively abstracting the detailed substructures. Chen et al.²⁵ proposed a chemistry-motivated graph neural network that uses templates to describe the net changes in electron configuration between the reactants and products. Recent advancements, such as the use of Modern Hopfield Networks for template prediction, have improved the performance of template-based CASP tools²⁶. Modern Hopfield Networks allow for the efficient prioritization of reaction templates by leveraging associative memory mechanisms, significantly improving the speed and reliability of retrosynthetic planning. Nevertheless, template-based methods often struggle with the rigidity of predefined rules, analyzed by Choe et al.²⁷, which can limit their ability to propose novel synthetic routes or adapt to emerging reaction types.

Computational retrosynthesis has traditionally been treated as a tree search problem, where each step involves searching for chemically feasible precursors to derive the product molecule. Early works have used greedy search⁴, but more recent efforts have adopted Monte Carlo tree search²⁸ and A*-like algorithms^{29,30}. These search methodologies attempt to optimize the exploration of synthetic routes by balancing the exploitation of high-confidence reaction pathways with the exploration of novel or underutilized transformations. However, their effectiveness is contingent on the quality and comprehensiveness of the reaction databases they utilize.

In addition to the different components of CASP, the integration of biocatalysis in computational retrosynthetic planning has gained attention for its potential to enhance sustainability. Biocatalysis leverages enzymatic transformations to achieve highly selective reactions under mild conditions, reducing the environmental footprint of chemical synthesis³¹. Incorporating enzymatic steps into synthetic pathways is particularly advantageous for pharmaceutical and fine chemical industries, where stringent purity and selectivity requirements are critical^{32,33}. Tools like RetroBioCat³¹ have demonstrated the ability to design biocatalytic cascades that complement synthetic routes by introducing enzymatic steps that are often more selective and environmentally friendly. However, even these specialized tools are limited by the availability of comprehensive enzymatic reaction data and their focus on narrow reaction types.

In this work, we introduce MHNpath, an ML-driven retrosynthetic tool designed to help chemists explore greener, cost-effective, and moderate-temperature synthetic pathways. First, we develop a robust template-prioritization model based on Modern Hopfield Networks³⁴, enhanced with Xavier initialization, dropout, and L2 regularization to improve stability and generalization. We additionally introduce two new evaluation metrics that enable a more comprehensive comparison of template prioritizers. Second, we implement a user-tunable scoring system that integrates practical constraints and include precursor cost, reaction temperature, and solvent toxicity, allowing chemists to steer the search toward more sustainable and operationally feasible routes. Third, we demonstrate MHNpath's effectiveness through case studies from PaRoutes³⁵ and ChemByDesign³⁶, showcasing its ability to identify viable, interpretable multi-step pathways. Together with its curated enzymatic and synthetic template datasets and a global greedy tree-search algorithm, MHNpath provides a flexible and data-efficient framework that addresses the limitations of existing CASP tools, particularly their reliance on extensive training and rigid template-based systems. The following sections will detail the methodology behind MHNpath, including the data processing pipeline (*Section 2.1*), the implementation of the Modern Hopfield Network-based template prioritizer (*Section 2.2*), and the tree search methodology and custom scoring employed for retrosynthetic exploration (*Sections 2.3 & 2.4*), as outlined in Figure 1. Furthermore, we will present results (*Section 3*) comparing MHNpath against existing CASP tools and provide case studies demonstrating its practical applications.



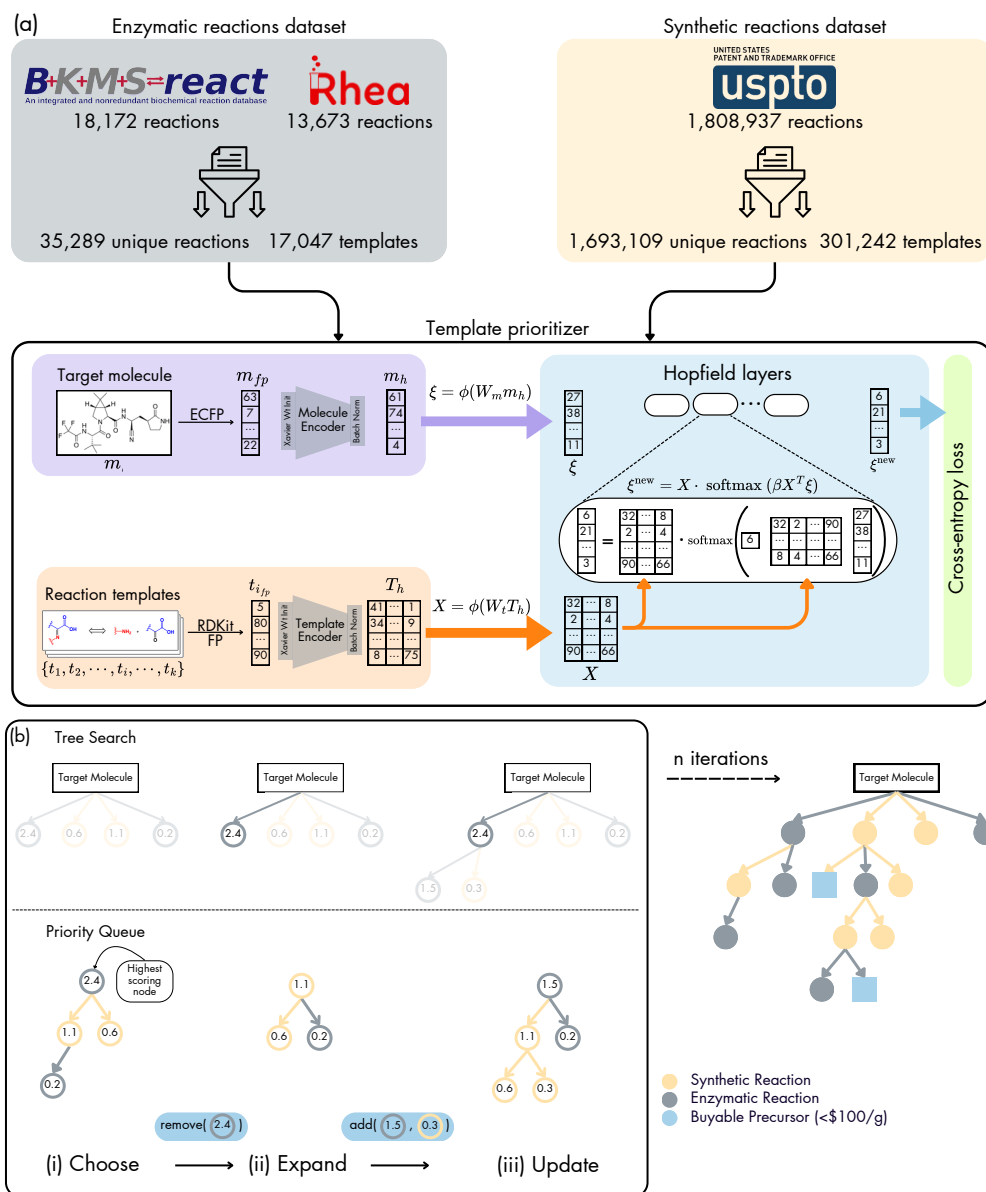


Figure 1: **(a) Data processing and model architecture.** The model development follows three primary stages. Step 1 (data curation): Raw enzymatic (BKMS-react, Rhea) and synthetic (USPTO) databases are cleaned and filtered to extract unique reaction templates. Step 2 (encoding): The target molecule m and the extracted reaction templates t are converted into ECFP and RDKit fingerprints, respectively, and passed through dedicated neural network encoders employing Xavier initialization and batch normalization to produce dense, continuous representations m_h and T_h . Step 3 (template prioritization): A Modern Hopfield Network projects m_h and T_h into a shared associative memory space, where an attention-like update mechanism produces the new molecule state ξ^{new} by retrieving the most applicable reaction templates, with the model optimized via cross-entropy loss. **(b) Tree search methodology.** A weighted global greedy tree search explores the retrosynthetic space iteratively in Step 4 (search and scoring). At each iteration, the algorithm (i) selects the highest-scoring node from a priority queue, (ii) expands it by applying the prioritized templates to generate new precursor nodes (grey circles for enzymatic, yellow for synthetic), and (iii) updates the priority queue using a multi-objective scoring function $P' = f(C', T', S')$ that promotes low precursor cost, green solvent usage, and moderate reaction temperatures. This cycle continues until buyable precursors (light blue squares, <\$100/g) are found, the maximum search depth is reached, or the allotted time limit is exhausted.

2 Methods

The MHNpath framework is presented in Figure 1. A modified Modern Hopfield Network³⁴ based architecture by Seidl et al.²⁶ is adapted. Each prioritizer (also referred to as model) accepts a target molecule as input and predicts a



ranked list of the most applicable reaction templates, which are applied to the target molecule to get precursors (nodes). The individual scores for precursor cost, reaction temperature, and solvent toxicity are calculated and summed using user-tunable weights. The precursors are further explored in a recursive manner using a global greedy tree search algorithm where the best-scoring unexplored node from the entire tree is explored first. Figure 1(b) gives an example of the multi-modal-guided global greedy tree search approach.

2.1 Dataset processing

Our study leveraged two comprehensive training datasets to develop the predictive component of the MHNpath framework. Additionally, specific external datasets were curated for benchmarking and evaluation tasks. In this context, we selected training data that provided a comprehensive representation of both enzymatic and synthetic reactions. The training dataset of enzymatic reactions was compiled from the BKMS³⁷ dataset and the Rhea³⁸ database, initially containing about 68,000 reactions. Following a rigorous preprocessing procedure that involved removing duplicates, invalid SMILES representations, and unbalanced reactions, we obtained 35,289 unique and valid reaction SMILES. From this refined enzymatic subset, we extracted 17,047 reaction templates using the RDKit³⁹ library. This subset was used to train a single *enzymatic template prioritizer*, utilizing an 80:10:10 random split for training, validation, and testing, respectively.

We also used a dataset of synthetic organic reactions for training which was sourced from the USPTO⁴⁰ dataset. Starting with 1,808,937 reactions and applying a series of preprocessing steps resulted in 1,693,109 refined reactions, from which 301,242 reaction templates were extracted in the synthetic subset. To manage computational resources and memory overheads due to the large number of data points in the synthetic subset, we employed a scaffold split technique using Datamol⁴¹ to ensure high chemical diversity while dividing the subset further into five equally sized subsets. This partitioning strategy allowed us to train five distinct *synthetic template prioritizers* (one per divided synthetic subset). Each divided subset underwent an 80:10:10 random split for model development. During inference, the predictions from these five models are aggregated using an ensemble approach. This strategy significantly enhances robustness with a single *synthetic template prioritizer* achieving a top-1 accuracy of 36.6% on its own test set and the ensemble method reaching a top-1 accuracy of 42.2% on the combined test set. In this context, top-1 accuracy denotes the model's ability to accurately prioritize the correct ground-truth template as the number one rank in the prediction list.

Once the predictive components of MHNpath were developed, we rigorously assessed the framework's performance in Sections 3.2 and 3.3 using three tasks as described below. These tasks required curating distinct datasets from the literature and provide unbiased performance metrics for systems not present in the training set.

- **PaRoutes benchmark:** To evaluate pathway reconstruction against patent literature, we utilized the PaRoutes³⁵ dataset. We applied a scaffold split to identify the ten most commonly occurring scaffolds within PaRoutes. From the clustering analysis, we selected 130 diverse target molecules to test our framework's ability to replicate and improve upon patent-derived "gold-standard" routes.
- **Novelty assessment (ChemByDesign):** To test generalization to unseen data, we selected 5 target molecules from ChemByDesign³⁶. We specifically filtered for pathways published after 2021 with fewer than six steps. This filtering ensured that the ground-truth reactions were not present in our USPTO training data (which was last updated in September 2016), providing a test of zero-shot pathway prediction.
- **Hybrid comparative set:** To benchmark against existing hybrid planners, we selected specific complex targets (e.g., dronabinol, arformoterol, 4-ethenyl-2-fluorophenol) used in prior studies by Levin et al.⁴ and RetroBioCat³¹. These molecules were manually curated to allow for a direct qualitative comparison of pathway length, cost, and enzyme usage.

2.2 Model development

We utilize the Modern Hopfield Network³⁴ based template prioritization architecture for our study. This architecture, initially introduced by Seidl et al.²⁶, consists of three main components: a molecule encoder, a reaction template encoder, and one or more stacked or parallel Hopfield layers.

The molecule encoder function, denoted as $h_w^m(\cdot)$, learns a relevant representation for the input molecule m . We utilize a fingerprint-based approach, specifically the extended connectivity fingerprint (ECFP), coupled with a fully connected neural network with weights w . This encoder maps a molecule to a dense representation $\mathbf{m}_h = h_w^m(m)$ of dimension d_m .

Similarly, the reaction template encoder function, denoted as $h_v^t(\cdot)$, learns relevant representations of reaction templates. We employ a fully connected neural network with RDKit template fingerprints as input. The function is applied to the



set of all templates \mathcal{T} , and the resulting vectors are concatenated column-wise into a matrix $\mathbf{T}_h = h_v^t(\mathcal{T})$ with shape (d_t, K) , where K is the number of templates. Both encoders utilize Xavier Initialization⁴² and Batch Normalization⁴³ to improve convergence and prevent vanishing gradients.

The core of the model consists of Hopfield layers, denoted as $g(\cdot, \cdot)$, which associate the molecule with the memory of templates. To perform the retrieval, the encoded molecule \mathbf{m}_h and the template matrix \mathbf{T}_h are projected into a common associative space to form the state pattern $\boldsymbol{\xi}$ and the stored patterns \mathbf{X} , respectively. The Hopfield layer then updates the molecule representation via an attention-like mechanism. The update rule yielding the new state $\boldsymbol{\xi}_{\text{new}}$ is defined as:

$$\boldsymbol{\xi}_{\text{new}} = \mathbf{X}\mathbf{p} = \mathbf{X} \cdot \text{softmax}(\beta\mathbf{X}^\top\boldsymbol{\xi}) \quad (1)$$

where β is a learnable scaling parameter (inverse temperature), \mathbf{p} is the vector of associations (probabilities) over the templates, and softmax is applied column-wise.

For the loss function, given a training pair (m, t) and the set of all templates \mathcal{T} , the model aims to maximize the probability assigned to the correct template t . We employ the negative log-likelihood as the loss function, optimizing parameters using stochastic gradient descent via the AdamW optimizer. To prevent overfitting, we apply dropout and L2 regularization. Additionally, a post-processing fingerprint-based substructure screen is used during inference to filter out chemically non-applicable templates.

We train two sets of models on our datasets: the *synthetic template prioritizer* (consisting of five models) and the *enzymatic template prioritizer*. We utilized PyTorch⁴⁴ version 1.9.0 and Python version 3.8 to implement the model. These models take a target molecule as input and rank the most applicable rules within the collected dataset. Given the same inputs, the outputs from the five *synthetic template prioritizers* are taken individually, and a final ranking of the templates is collated using the highest predicted score. Furthermore, to fine-tune the models for optimal performance, hyperparameter tuning was performed using a one-factor-at-a-time (OFAT) approach. We tuned number of epochs, concatenation threshold, dropout rate, learning rate and some hopfield parameters, more information on the parameters and the values chosen can be found in Supplementary Sections 1–4. Each experiment involved modifying a single hyperparameter from a baseline configuration while keeping all other settings constant. The performance of each model was assessed using an evaluation score, with lower validation loss indicating better performance. Training and validation loss curves and accuracy metrics for the enzymatic and synthetic template prioritizers are provided in Supplementary Figures 1 and 2.

2.3 Searching algorithm

We employ a weighted global greedy tree search approach to explore synthesis pathways for a target molecule. This method leverages the template prioritizer models to identify the most applicable reaction rules. Specifically, the MHN generates a probability distribution over the predefined reaction templates, and we apply a cutoff by selecting only the top- k (also user-tunable) most confident applicable rules. These filtered rules are then applied to the target molecule to derive precursor molecules. In this search framework, precursor molecules are represented as nodes in a tree, while reactions and their associated conditions serve as edges. However, not all precursors may be readily available or affordable. We define “buyable” as being available for purchase at a cost of under \$100/g. A representative synthetic pathway tree illustrating the search output for 2-phenoxyethanamine is provided in Supplementary Figure 3. Algorithm 1 outlines the weighted global greedy tree search algorithm. Excluding API latency, the runtime complexity is $\mathcal{O}(b^d \log V)$, where b is the effective branching factor bounded by the user-tunable top- k parameter, d is the maximum search depth, and $\log V$ reflects the cost of priority queue operations over V nodes. The search algorithm proceeds as follows:

- **Initialization:** We begin by initializing the search with the target molecule as the *start_node*. This node is added to a priority queue with high priority.
- **Global greedy search:** The main search loop iterates until the priority queue is empty. At each iteration, the node with the highest priority is popped from the queue.
- **Goal check:** If the node meets the goal criteria (e.g., low cost or maximum depth), the algorithm continues to the next iteration.
- **Rule application:** For the current node, we retrieve the top- k transformation rules ranked by the MHN’s predicted probabilities. We apply each valid rule to generate new chemical structures. The physical and operational properties of these new structures (such as cost, temperature, and solvent score) are then calculated.
- **Node insertion:** Each new structure is encapsulated in a *new_node*, which is added to the current node’s subtrees and inserted into the priority queue based on its multi-objective score.



- **Scoring:** Nodes are strictly prioritized in the queue based on a user-tunable criterion $P' = f(C', T', S')$ that includes precursor cost, reaction temperature, and solvent toxicity. The highest-scoring nodes are explored first.
- **Termination:** The search continues iteratively until buyable precursors are found, the tree is fully explored up to a specified maximum depth, or the allotted search time is exhausted.

Algorithm 1 Weighted Global Greedy Tree Search for Chemical Pathways

```

1: Initialize
2:  $N_0 \leftarrow$  initial chemical structure with cost  $C_0$ 
3:  $PQ \leftarrow \{\}$  ▷ Priority queue
4: add  $N_0$  to  $PQ$  with score  $P_0 = \infty$ 
5: while  $PQ \neq \emptyset$  do
6:    $N \leftarrow$  pop node with highest score from  $PQ$ 
7:   if goal_check(N) then
8:     continue
9:    $R \leftarrow$  find_top_k_rules(N, k) ▷ Filter by MHN confidence
10:  for  $r \in R$  do
11:     $N' \leftarrow$  apply_rule(r, N)
12:     $C' \leftarrow$  calculate_cost(N')
13:     $T' \leftarrow$  calculate_temp(N')
14:     $S' \leftarrow$  calculate_solvent_score(N')
15:     $P' \leftarrow$  calculate_score(C', T', S')
16:    create new node  $N'$  with  $C', T', S'$ 
17:    add  $N'$  to subtrees of  $N$ 
18:    insert  $N'$  into  $PQ$  with priority  $P'$ 
19: function FIND_TOP_K_RULES(N, k)
20:   return top  $k$  applicable rules for  $N$  based on MHN probability
21: function APPLY_RULE(r, N)
22:   return new chemical structure from applying  $r$  to  $N$ 
23: function CALCULATE_SCORE(C', T', S')
24:    $P' \leftarrow f(C', T', S')$  ▷ Compute multi-objective score
25:   return  $P'$ 
26: function GOAL_CHECK(N)
27:   return True if  $N$  meets goal conditions (low  $C$  or max depth), else False
28: Output
29: Print resulting tree of chemical reactions
  
```

2.4 Scoring system

We introduce a user-tunable scoring methodology to evaluate synthesis pathways. In this approach, individual scores for three features are obtained and summed based on user-tunable weights, allowing users to prioritize features according to their preferences during the pathway search. A detailed workflow diagram of our inference and scoring workflow is shown in Figure 2, and representative ablation studies illustrating the effect of the user-tunable scoring weights on route selection are presented in Supplementary Figures 11–13. The first feature of the score is the cost. We utilize the ASKCOS⁴⁵ buyable dataset, Molport⁴⁶, Mcule⁴⁷, and Chemspace⁴⁸ APIs with the ChemPrice⁴⁹ library to determine the cost and availability of molecules. A molecule is considered buyable if it can be purchased for less than \$100/g^{4,50}. Users have the flexibility to adjust this threshold based on their objectives. The score is calculated as $-\frac{\text{price}}{500}$, as we aim to maximize the score, thus exploring reactions involving cheaper precursors first. The value 500 is chosen as a normalizer because we define anything costing above \$500/g as non-buyable; the user can modify this threshold. These threshold align with the affordability criteria for research-grade building blocks utilized in major chemical databases⁵⁰.

The second feature of the score is the reaction temperature. We employ a CASP tool developed by Gao et al.⁵¹ to predict the temperature at which a reaction might occur. This prediction is essential for exploring pathways where reactions may not have been previously documented or experimented on. To find the predicted temperature, we take the weighted average of the top ten predictions made by the tool, a strategy suggested by the authors to improve accuracy. The score is calculated as $-\frac{\text{temperature}}{300}$, as we aim to maximize the score, thus exploring reactions involving lower temperatures



first. The value 300 is chosen as a normalizer because we define any reaction requiring over 300° C as non-practical; this cutoff reflects the upper operating limits of standard synthetic laboratory equipment⁵².

The third feature of the score is the solvent and reagent greenness-toxicity score. We use the same CASP tool developed by Gao et al.⁵¹ to predict the solvent and reagent required for a reaction. We curate a toxicity and greenness dataset, assigning a score of -1 for toxic molecules, 0 for neutral molecules, and +1 for green or natural molecules. We use the ACS Solvent Selection Guide⁵³ to classify 100 commonly used solvents. We also use the SuperNatural 3.0 dataset⁵⁴, containing 350,000 natural products, and the T3DB dataset⁵⁵ to classify 4,000 toxic molecules. This ternary scoring system serves as a heuristic to steer the tree search toward environmentally favorable regions of chemical space. While this approach does not currently account for process-specific metrics such as atom economy, E-factors, or reagent concentration, it provides a computationally efficient means of penalizing more hazardous precursors. This modular design allows the $-1/0/+1$ baseline to be substituted with more granular, continuous toxicity models as they become available. Furthermore, because these scores are additive, users can adjust the global weight of the greenness feature to balance environmental safety against precursor cost and temperature constraints.

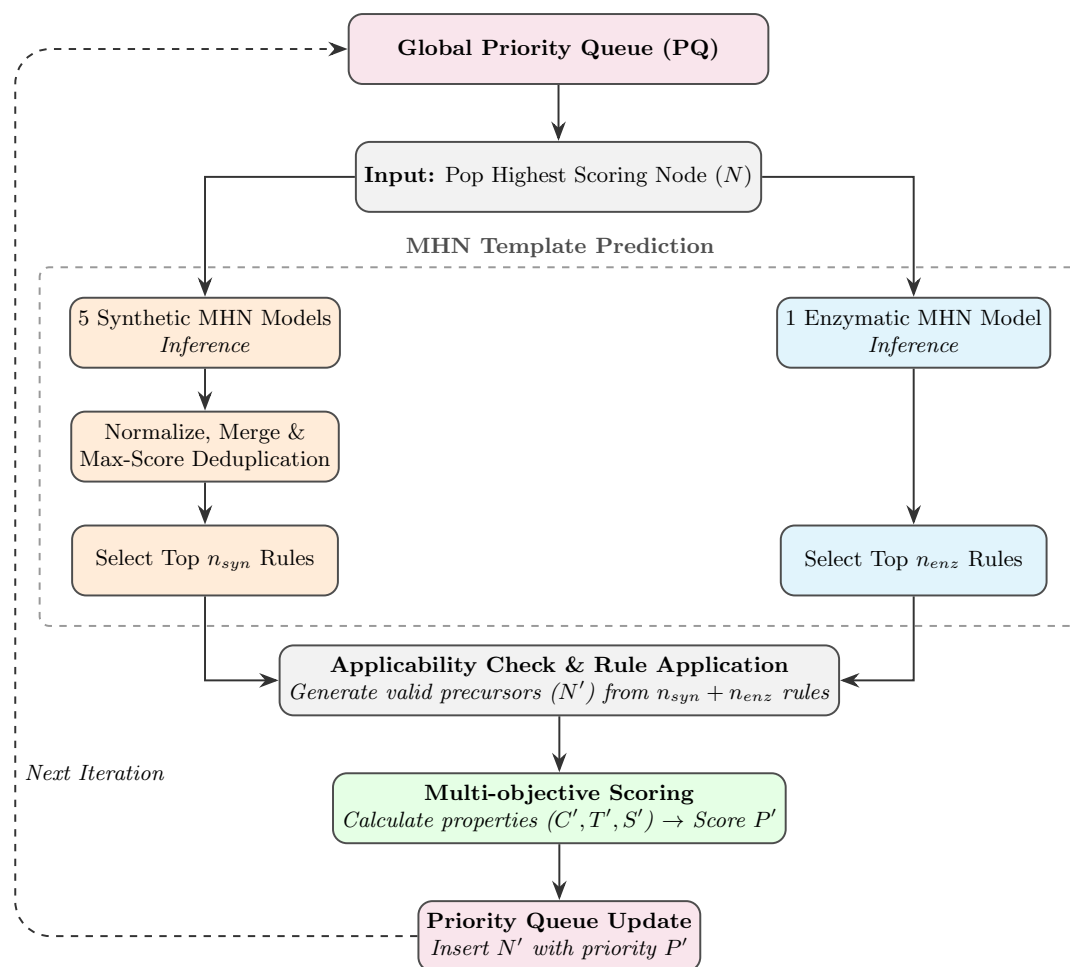


Figure 2: **Inference and Scoring Workflow.** This schematic illustrates the step-by-step pipeline executed during each iteration of the weighted global greedy tree search (Algorithm 1). First, the highest-scoring target node (N) is extracted from the global priority queue. The node is concurrently passed to the synthetic and enzymatic template prioritizers. For the synthetic pathway, an ensemble of 5 MHN models predicts applicable rules. These ranked lists are normalized, merged, and deduplicated (retaining the highest score per unique rule) to isolate the top n_{syn} templates. Simultaneously, the single enzymatic MHN model yields the top n_{enz} templates. During the *Applicability Check*, this combined set of $n_{syn} + n_{enz}$ rules is applied to N to generate chemically valid precursor nodes (N'). Each valid precursor undergoes *Multi-objective Scoring* to calculate its specific cost (C'), temperature (T'), and solvent (S') constraints, culminating in a final priority score (P'). Finally, a *Priority Queue Update* inserts these newly scored precursors into the global queue, directing the next search iteration toward the most highly optimized pathways.



3 Results and Discussion

In this section, we evaluate the performance of the MHNpath framework in prioritizing reaction templates and generating feasible retrosynthetic pathways. We benchmark our predictive models against established baselines and analyze the practical utility of the proposed scoring system in guiding pathway discovery.

3.1 Model performance

Model	Lit. Rule Accuracy				Avg. Applicable Rules				Presence of Applicable Rule			
	T1	T10	T50	T100	T1	T10	T50	T100	T1	T10	T50	T100
DNN ⁴	0.100	0.312	0.442	0.485	0.128	0.898	4.876	14.005	0.128	0.539	0.900	0.973
MHNreact ²⁶	0.181	0.566	0.760	0.809	0.199	1.004	5.647	16.388	0.199	0.531	0.861	0.967
Ours (Enz)	0.183	0.576	0.763	0.814	0.201	1.076	6.125	17.380	0.201	0.553	0.884	0.976
Ours (Syn) [†]	0.422	0.776	0.899	0.927	0.891	6.315	19.703	30.739	0.891	0.968	0.994	0.996

Table 1: **Performance Metrics for Template Prioritization.** The best, second best, and third best results for the enzymatic dataset are color coded. The table is divided into three sets of metrics. The first set (columns 2-5) shows the accuracy of the presence of the literature rule in the top predictions (T1, T10, T50, T100). The second set (columns 6-9) represents the average number of applicable rules in the top predictions. The third set (columns 10-13) indicates the accuracy of the presence of at least one applicable rule in the top predictions.

[†] The synthetic model is trained and tested on a separate, larger dataset, while the other three models are trained and tested on the same dataset.

To evaluate the effectiveness of our synthesis planning framework, we employed a comprehensive set of metrics designed to assess both the accuracy and applicability of the predicted reaction templates. Table 1 presents the performance of our model in comparison to existing state-of-the-art models, using three primary metrics:

- Accuracy of the presence of the literature rule in the top n predictions (T1, T10, T50, T100):** This metric evaluates the model's ability to prioritize reaction templates documented in the literature. The "literature rule" refers to the reaction template extracted from the corresponding reaction in our dataset. This metric is crucial for assessing how well the model replicates known synthetic pathways. This is the commonly used metric in recent works^{4,26}. While this metric is important for evaluating the replication of known synthetic pathways, assessing the model's ability to propose novel and feasible pathways is equally critical. Therefore, we also consider the following metrics, which provide a more comprehensive view of the model's performance.
- Average number of applicable rules in the top n predictions:** This metric provides insight into the diversity and feasibility of the proposed reaction templates. A rule is considered "applicable" when the RDKit library successfully applies the transformation to the target molecule, resulting in valid precursor structures. This metric is particularly important as it indicates the model's capability to explore broader chemical space and suggest alternative synthetic routes that may not be present in the literature but are chemically plausible. This is critical for our tree search, as a higher density of valid rules in the top predictions (T50, T100) expands the branching factor of the search tree with high-quality and newer candidates.
- Accuracy of the presence of at least one applicable rule in the top n predictions:** This metric complements the other two by measuring the model's practical utility in retrosynthetic analysis. It assesses how often the model suggests at least one viable synthetic step, ensuring the progression of retrosynthetic analysis in chemically plausible directions. High scores here ensure the search algorithm rarely hits "dead ends".

The introduction of these additional metrics (2 and 3) addresses the limitations of traditional evaluation methods, which often focus solely on predicting literature rules. While the accuracy of predicting literature rules is important, it does not fully capture the model's ability to suggest novel pathways. By incorporating metrics that consider rule applicability and diversity, we provide a more comprehensive evaluation that aligns with the exploratory nature of retrosynthetic planning.

Our enzymatic model (Ours (Enz)) demonstrates good performance across most metrics compared to the baseline DNN (Deep Neural Network)⁴ and MHN²⁶ models. It is important to contextualize the baseline performance, the DNN achieves a Top-1 accuracy of only 10%, which underscores the inherent complexity of the retrosynthetic prediction task. Unlike standard classification problems, template prioritization involves selecting the correct chemical transformation



from a massive search space of valid reaction templates. In this high-dimensional context, low absolute accuracy scores are standard, and incremental gains represent significant practical improvements.

First, compared to the standard DNN baseline⁴, our model delivers a large improvement. The DNN achieves a Top-1 literature accuracy of only 10% and a Top-1 applicability presence of 12.8%. Our model nearly doubles these figures to 18.3% and 20.1%, respectively. Given the high-dimensional search space (selecting 1 out of >17,000 templates), this jump represents a significant improvement in the model's ability to prioritize relevant chemistry. This trend is mirrored in the applicability metrics, indicating that our model not only retrieves ground truth better but also generates a higher number of valid chemical precursors.

Second, while our architecture shares similarities with the MHN model²⁶, our specific enhancements (Xavier initialization, rigorous dropout, and hyperparameter tuning) yield better convergence and robustness, particularly when predicting precursors for structurally complex or challenging target molecules. For example, in the *Avg. Number of Applicable Rules* metric at T50, our model outperforms the MHN baseline (6.125 vs. 5.647). This indicates that our model retrieves approximately 8.5% more valid chemical options in the top-50 candidates. This increased density of applicable rules provides the global greedy search algorithm with a richer pool of precursors, reducing the likelihood of missing a viable pathway.

Finally, the synthetic ensemble model (Ours (Syn)) demonstrates exceptional performance on the USPTO dataset, achieving a T1 accuracy of 42.2%. Most notably, it exhibits a massive applicability rate: the Top-10 predictions yield, on average, 6.315 valid reaction rules. This high applicability ensures that the synthetic branch of our hybrid planner can almost always identify multiple feasible precursors, allowing the scoring system to aggressively optimize for cost and environmental impact without running out of chemical options.

It is important to note that the baseline DNN⁴, the MHN model²⁶, and our enzymatic model (Ours (Enz)) were trained and evaluated on the exact same enzymatic dataset splits, ensuring a direct and fair comparison.

3.2 Comparison with pathways from the literature

We assessed our synthesis planning framework by benchmarking it against established datasets and frameworks, including PaRoutes³⁵ and ChemByDesign³⁶. Precisely, we assess the pathway lengths generated by our framework and report metrics on the number of replicated pathways and the average number of predicted pathways per molecule. We also present some representative examples involving a hybrid approach combining enzymatic and synthetic pathways in Figure 3(a) and Supplementary Figures 4 and 5. This comparative analysis aimed to assess our framework's accuracy, efficiency, and versatility in predicting viable synthetic routes for complex target molecules.

PaRoutes³⁵ is a robust benchmarking framework for evaluating multi-step retrosynthesis methods. It comprises two datasets of 10,000 synthetic routes derived from the patent literature alongside a curated list of purchasable molecules and reactions suitable for training retrosynthesis models. For this study, we utilized a scaffold split technique to select pathways associated with the ten most commonly occurring scaffolds in the PaRoutes dataset, which are given in Supplementary Section 5. This approach enabled us to efficiently assess our framework's ability to predict synthetic routes that align with those documented in the patent literature while ensuring chemical diversity.

ChemByDesign³⁶, on the other hand, is an online platform that organizes experimentally verified reaction pathways by name, year, and author. To ensure an unbiased evaluation, we focused on pathways discovered after 2021 that are under six steps long. These pathways were intentionally excluded from our training dataset to test the predictive capabilities of our framework on novel and unexplored synthetic routes. This strategy provided a unique opportunity to evaluate how effectively our framework generalizes to unseen data.

The results of this comparative analysis are summarized in Figure 3. Figure 3(a) presents a representative tree of reaction pathways generated by our framework for a target molecule in the PaRoutes dataset. This evaluation was conducted on an expanded test set comprising 130 molecules from the PaRoutes dataset and 5 novel targets from ChemByDesign. Notably, our predicted pathway utilizes inexpensive precursors and environment-friendly solvents, such as ethanol and methanol, demonstrating its potential for sustainable synthesis planning. The tree also highlights alternative routes that leverage naturally occurring molecules while avoiding toxic solvents like dichloromethane (DCM); further examples of these green alternatives are illustrated in Supplementary Figures 7–9. These results underscore the framework's ability to prioritize green chemistry principles without compromising synthetic feasibility and cost.



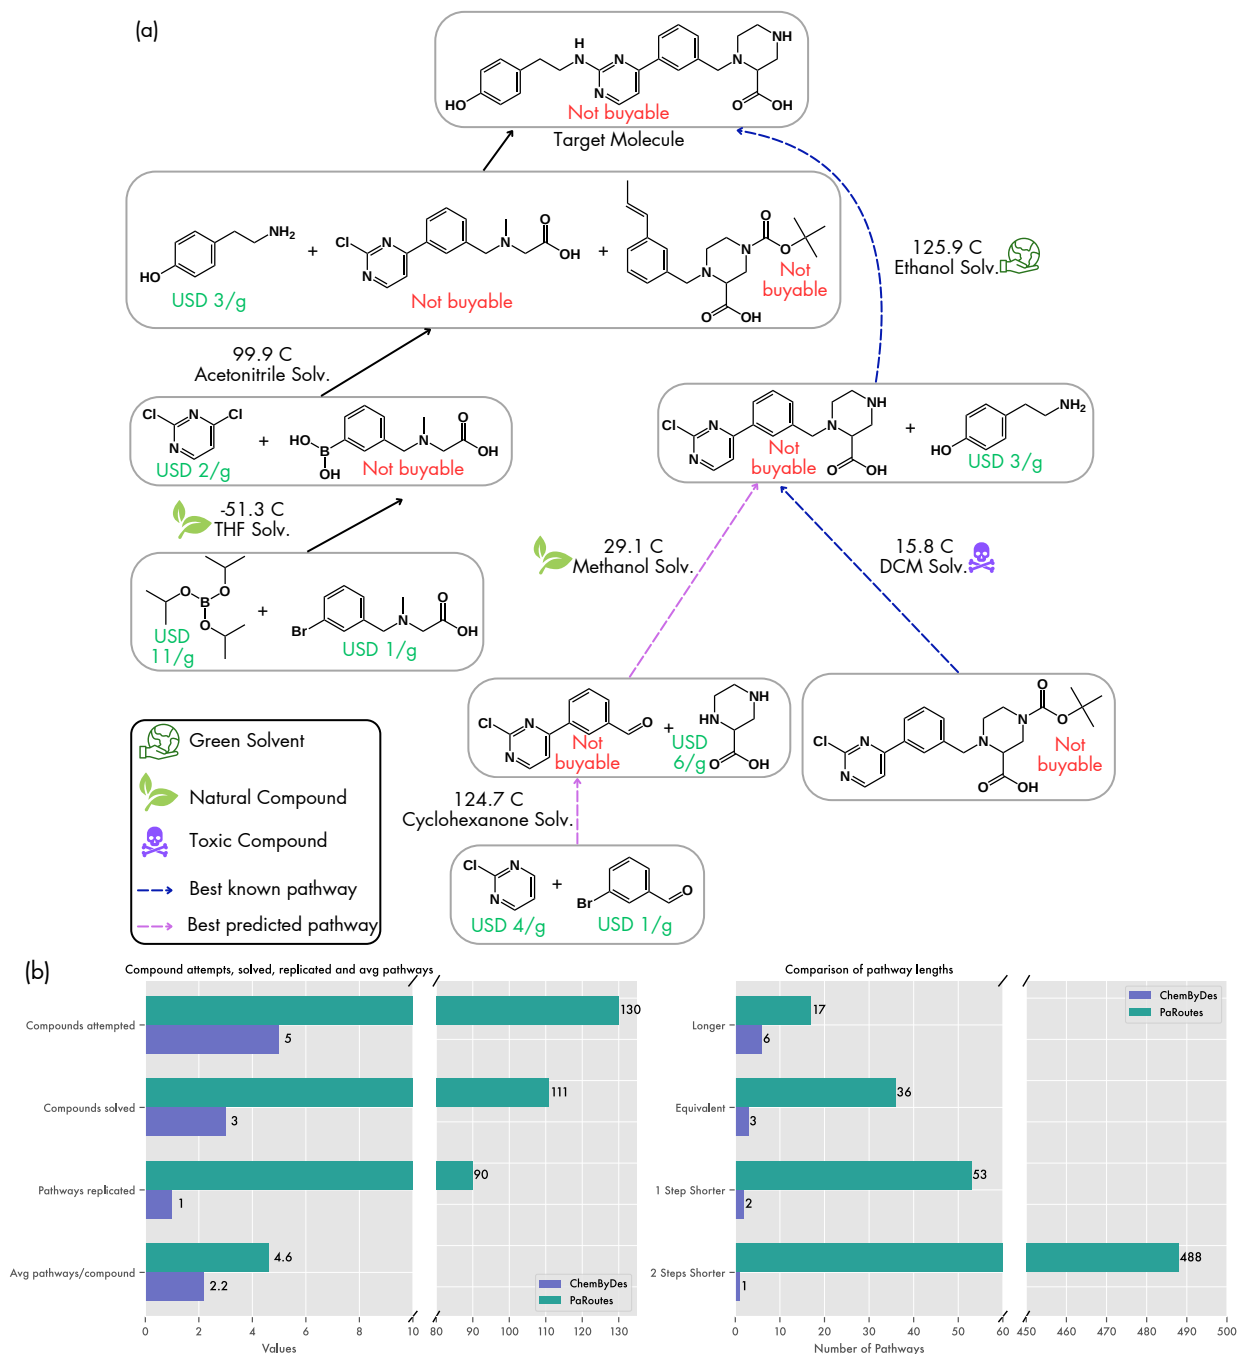


Figure 3: (a) **Tree of reaction pathways.** The tree shows a representative example for a pathway presented in PaRoutes³⁵. Our predicted pathway is producible using cheap precursors and less toxic, naturally occurring solvents. (b) **Performance metrics for literature comparison.** These plots present the number of molecules solved, the average number of pathways predicted, and the distribution of predicted pathway lengths.



Figure 3(b) provides quantitative performance metrics for our framework compared to PaRoutes and ChemByDesign. The first bar plot illustrates the number of molecules attempted and successfully solved by each method. Our framework solved 114 out of 135 molecules attempted. A representative example of a successfully replicated PaRoutes pathway, annotated with reaction conditions and precursor costs, is illustrated in Supplementary Figure 7. Furthermore, our framework replicated 91 known pathways from the literature while suggesting a high number of pathways on average per molecule (4.6 and 2.2). These results demonstrate its ability to explore diverse chemical spaces and propose multiple viable options for synthetic planning.

Finally, the third bar plot compares the lengths of predicted pathways relative to known literature routes. Our framework identified 55 shorter or 39 equivalent-length pathways, with 488 cases yielding routes two steps shorter than those documented in PaRoutes, including two notable molecules N-(3-[5-Chloro-2-(difluoromethoxy)phenyl]-1-2-[(4-pyridinylmethyl)amino]ethyl-1H-pyrazol-4-yl)pyrazolo[1,5-a]pyrimidine-3-carboxamide and 4-(1-tert-butyl-4-oxo-5H-pyrazolo[4,3-c]pyridin-3-yl)thiophene-2-carboxamide, and one molecule (Lupinine) yielding a route two steps shorter than that documented in ChemByDesign. This capability to optimize pathway length is particularly valuable in industrial settings where shorter synthetic routes can lead to significant cost savings and improved process efficiency.

Supplementary Figure 9 illustrates an alternative pathway for Lupinine⁵⁶ (found in ChemByDesign) synthesis predicted by MHNpath. Our framework identified a streamlined three-step route with moderate reaction conditions (12.58°C, 11.1°C, and -23.81°C) and low precursor costs (\$87.36/g and \$0.10/g), contrasting with the five-step approach by Wang et al. that spans temperatures from -78°C to 85°C. The test set of utilized from PaRoutes has been provided in the Supplementary Data.

3.3 Comparison with other models

We benchmark the performance of MHNpath against two retrosynthetic planning tools: RetroBioCat³¹ and the hybrid enzymatic-synthetic planner developed by Levin et al.⁴, as illustrated in Figure 4. Precisely, we assess the pathway lengths generated by our framework and report metrics on the number of replicated pathways and the average number of predicted pathways per compound. We also present some case studies involving a hybrid approach combining enzymatic and synthetic pathways in Figure 4(i) and Supplementary Figure 6. These comparisons highlight the advantages of our framework in terms of pathway length, cost-effectiveness, and the ability to replicate and improve upon previously reported pathways.

Although synthetic planning tools such as AiZynthFinder⁵⁷, ASKCOS⁴⁵, Retro*²⁹, and MHNreact²⁶ represent significant benchmarks in the field, this section prioritizes comparisons with hybrid planners. The predictive capabilities of the underlying MLP and MHN-based engines used by these synthetic tools have already been quantitatively evaluated in Section 3.1 (Table 1). We exclude a full tree search comparison with these platforms as they are restricted to purely synthetic routes and lack the integrated multi-objective scoring criteria central to MHNpath, specifically real-time cost, toxicity, and temperature optimization. Furthermore, we do not report search time metrics because our framework relies on live API calls to retrieve dynamic pricing information. Since the total search duration is dominated by network latency rather than algorithmic efficiency, direct runtime comparisons with tools utilizing static building block datasets would be unrepresentative. Consequently, comparing global success rates between conventional planners, which optimize primarily for route brevity and model confidence, and MHNpath would be fundamentally misaligned, as our algorithm routinely bypasses traditionally 'successful' shortest paths in favor of longer, but demonstrably greener and more cost-effective alternatives.

RetroBioCat³¹ is a widely used platform for designing biocatalytic cascades. RetroBioCat facilitates the construction of selective and efficient biocatalytic pathways by leveraging an expanding enzyme toolbox and encoded reaction rules. Its strength lies in its ability to identify promising enzyme-specific routes, validated through several literature examples. However, as shown in Figure 4 (i), our framework finds shorter pathways than RetroBioCat for many molecules like 4-ethenyl-2-fluorophenol and 2-phenylpiperidine. Specifically, MHNpath successfully solved all five molecules attempted using RetroBioCat data, with an average of 4.2 pathways per molecule. The molecules attempted are listed in Supplementary Section 18. Moreover, our framework excelled in uncovering multiple shorter pathways and various pathways of equivalent length, offering users a range of options encompassing different enzymes, starting molecules, and temperature conditions.

Supplementary Figure 6 compares the pathways generated by RetroBioCat and MHNpath. In the top panel, while the RetroBioCat pathway for a fluorinated compound involves multiple enzymatic steps (TPL, TAL, and DC), our approach achieves the same transformation using a single enzymatic step with TPT at 99°C, thereby significantly reducing both the complexity and the precursor cost to \$6.03/g. In the bottom panel, for an amine compound, the RetroBioCat route requires a three-step enzymatic cascade (CAR, TA, and IRED) that depends on cofactors such as NADPH and ATP. In



contrast, our hybrid pathway utilizes only two enzymatic steps (PT at 54°C and PP at 10.1°C), with considerably lower precursor costs (\$0.1/g and \$1.8/g, respectively).



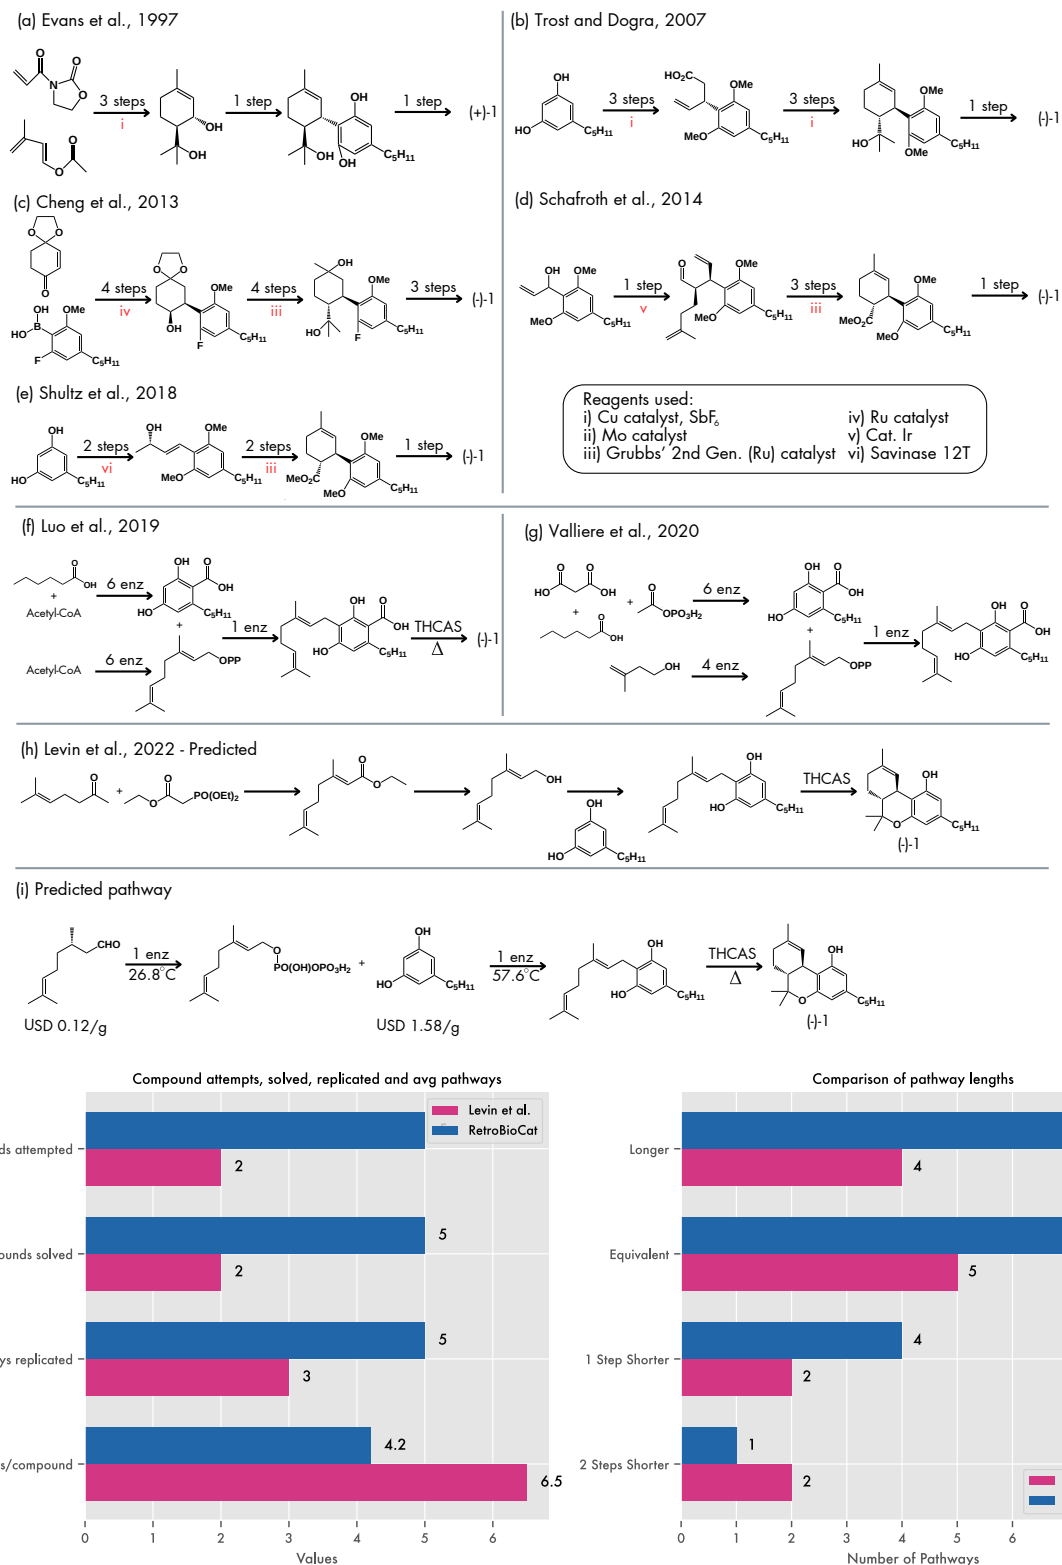


Figure 4: (a-g) **Published pathways.** Overview of existing pathways to produce dronabinol. (h) **Previously predicted pathways by Levin et al.**⁴. Four-step reaction pathway to produce dronabinol as predicted by Levin et al.⁴. (i) **Our predicted pathway.** Three-step reaction pathway to produce dronabinol from cheap precursors in ambient temperatures. We also replicated some of the other pathways. (j) **Performance metrics for comparison with other models.** These plots present the number of molecules solved, the average number of pathways predicted, and the distribution of predicted pathway lengths^{4,31}.



The hybrid enzymatic-synthetic planner by Levin et al.⁴ represents another state-of-the-art approach that combines neural networks trained on extensive reaction databases with enzymatic transformations. Levin's model offers hybrid strategies for complex molecules such as THC and R,R-formoterol, showcasing the potential of enzyme-synthetic integration. However, our results demonstrate that MHNpath not only replicates several of Levin's proposed pathways but also identifies novel alternatives that are shorter and more cost-effective.

As shown in Figure 4 (h) and (i), Levin's predicted pathway for dronabinol involves a four-step synthesis. In contrast, our framework proposes a three-step pathway that reduces synthesis costs to \$0.12/g at ambient temperatures. This significant cost reduction is achieved by incorporating an optimized enzymatic step that eliminates the need for high-temperature reactions, underscoring the practical advantages of our approach. Although the final step of the predicted pathways remains unchanged, our framework excels in exploring diverse routes to the penultimate molecule, owing to its global-greedy tree search algorithm and scoring methodology. In addition to replicating published pathways from Levin et al., our framework identifies four shorter additional pathways.

Furthermore, Supplementary Figure 8 presents a novel four-step synthetic pathway for the synthesis of arformoterol predicted by our model. It starts from inexpensive precursors (\$3.11/g, \$1.30/g, \$0.51/g, and \$0.24/g). In contrast, Levin et al.'s approach involves a more complex five-step biocatalytic cascade that requires multiple enzymes and cofactors. The detailed annotations of reaction conditions in our pathway, such as temperature, solvent, and reagent specifics, demonstrate the practical applicability of our hybrid strategy.

A key advantage of MHNpath is its ability to balance exploration and exploitation during retrosynthetic planning. This is evident from the higher average number of pathways generated per molecule (6.5 and 4.2), which provides chemists with a broader range of options for optimizing synthesis strategies based on reaction temperature, cost, or solvent toxicity.

4 Conclusions

The development and implementation of the MHNpath tool represent progress over existing retrosynthetic frameworks in computer-aided synthesis planning. By leveraging machine learning, we have created a framework that can efficiently predict retrosynthetic pathways and facilitate the exploration of diverse chemical spaces. Our template prioritizer outperformed existing methods, as demonstrated by its higher accuracies across specific metrics and increased number of applicable rules. Importantly, our framework was able to replicate gold-standard pathways from PaRoutes as well as novel experimental reaction routes reported in ChemByDesign. In some instances, the model not only reproduced known synthesis routes but also identified alternative pathways that were shorter and more cost-effective. For example, our approach discovered a three-step synthetic route for dronabinol, reducing the synthesis cost to \$0.12/g as compared to a previously reported four-step pathway. The framework also demonstrated partial retrosynthetic decomposition of more structurally complex targets, as illustrated in Supplementary Figure 10, though reaching fully buyable terminal states for such targets remains an open challenge.

MHNpath underscores how ML can be utilized in reducing the manual workload and minimizing the trial-and-error strategies traditionally associated with chemical synthesis. Combined with a tree search-based strategy, the machine learning model allows for the automated and efficient prediction of optimal reaction pathways. The tool's user-tunable criteria allow researchers to prioritize different aspects of the synthesis process, such as cost, reaction temperature, and toxicity of solvents and reagents. This adaptability is particularly beneficial for experimental chemists who can tailor the tool to meet their research needs and constraints.

Recent years have seen a surge in chemistry-focused large language models (LLMs) such as BatGPT-Chem⁵⁸ and ChemDFM⁵⁹, which bypass rigid template libraries to offer broad generalization across diverse chemical spaces. However, several properties of these models make them poorly suited as the core engine of a constraint-driven CASP tool. First, their non-deterministic, autoregressive generation processes introduce a non-negligible risk of producing chemically invalid intermediates or synthetically implausible steps, errors that propagate and compound across iterations in a deeply branched tree search. Second, state-of-the-art chemistry LLMs typically require on the order of 13 to 15 billion parameters, creating substantial computational bottlenecks when integrated into iterative search algorithms that may query the model hundreds of times per target molecule. Third, and most critically, it remains exceptionally difficult to algorithmically enforce hard physical constraints, such as strict precursor cost ceilings or solvent toxicity indices, through text-based prompting alone.

By contrast, the MHN-based template prioritization approach used in MHNpath explicitly maps target molecules and the full template library into a shared associative memory space, producing deterministic and structurally valid outputs at every retrosynthetic step by construction. Although MHNpath operates on a predefined template library, it is not strictly limited to historical precedent. The extraction of 301,242 localized transformation rules from approximately



1.7 million USPTO reactions allows the model to apply generic reaction templates to novel substrates, effectively generating novel multi-step pathways without sacrificing intermediate chemical validity. The multi-objective scoring function $P' = f(C', T', S')$ enforces physical constraints algorithmically and with full transparency, granting synthetic chemists a degree of precise operational control that generative LLMs currently struggle to guarantee.

Beyond architectural differences, MHNpath's utility is designed to function as a decision-support system rather than a black-box generator. While the multi-objective optimization identifies pathways that are computationally "greener" or more "cost-effective," these outputs serve as prioritized hypotheses for expert evaluation. By quantifying the trade-offs between temperature, cost, and solvent toxicity, the framework provides synthetic chemists with a structured "shortlist" of routes that align with specific laboratory constraints, significantly narrowing the vast search space before high-cost experimental validation is attempted.

Despite its capabilities, MHNpath has some limitations that outline future work directions. One potential issue is the reliance on predicted pathways without sufficient experimental validation, meaning that while the tool can suggest plausible synthetic routes, it cannot guarantee their success in practice and should be used as a starting point for further experimental investigation rather than a definitive solution. Additionally, the tool cannot effectively address enantiomer selectivity issues, as it does not inherently account for the stereoselective outcomes of reactions involving chiral molecules. Furthermore, MHNpath is limited in its ability to provide detailed mechanistic insights into the predicted reactions, a drawback for users who require a deeper understanding of the underlying processes to optimize and troubleshoot reactions. Its overall accuracy and reliability are also heavily dependent on the quality and comprehensiveness of the underlying datasets; incomplete or biased data can lead to inaccurate predictions and missed opportunities for novel synthetic routes. To address these challenges, future work will focus on integrating enantioselective predictions by incorporating enantiomer-specific data and developing algorithms capable of predicting stereoselective outcomes, providing more mechanistic information through the integration of mechanistic databases, and continuously expanding and diversifying the datasets via collaborations and feasibility tests with experimental chemists.

In summary, MHNpath demonstrates the potential of ML-driven tools in advancing CASP. Our results highlight this promise: for example, the enzymatic model achieved a top-1 accuracy of 18.3%, a marked improvement over the 10% of the DNN baseline and comparable to the 18.1% of MHN. On the other hand, the synthetic ensemble model reached impressive 42.2% top-1 accuracy. Moreover, our framework delivered an average of approximately 6 applicable rules in the top 10 predictions, highlighting its ability to explore diverse chemical spaces. In comparisons with pathways in the literature, MHNpath successfully solved 114 out of 135 molecules attempted and even identified a three-step synthetic route for dronabinol that reduces costs to \$0.12/g at ambient temperatures, outperforming a competing four-step pathway. However, addressing its limitations and mitigating potential misuse remains essential to ensure the long-term success and reliability of MHNpath for assisting in synthesis of complex organic molecules.

Data Availability

The synthetic dataset utilized in this study was obtained from the USPTO⁴⁰ repository, while the enzymatic dataset was sourced from the BKMS³⁷ and RHEA³⁸ databases. All datasets are publicly accessible and open-source. The processed datasets generated and analyzed during the study and the trained model weights are available on [Figshare](#).

Code Availability

The code used for data processing, model training, inference and the instructions to run our framework are available at <https://github.com/MSRG/mhnpath>.

Acknowledgement

This research has received funding from the research project: "Quantum Software Consortium: Exploring Distributed Quantum Solutions for Canada" (QSC). QSC is financed under the National Sciences and Engineering Research Council of Canada (NSERC) Alliance Consortia Quantum Grants #ALLRP587590-23. The authors also acknowledge the funding received from the University of Toronto's Data Sciences Institute (DSI) via its Catalyst Grant as well as the Canada Research Coordinating Committee's (CRCC) New Frontiers in Research Fund (NFRF) for their continued support. S. P. and V. K. P. thank DSI for providing financial support during the summer of 2025 via the Summer Undergraduate Data Science Research Opportunities Program and the Postdoctoral Fellowship, respectively. S. P. would also like to thank MolPort, Mcule and ChemSpace for providing API keys and access to a cost related database



of chemical compounds. V.K.P. is grateful for the computational resource support provided by the Digital Research Alliance of Canada.



References

- [1] Szymkuć, S. *et al.* Computer-assisted synthetic planning: The end of the beginning. *Angewandte Chemie International Edition* **55**, 5904–5937 (2016).
- [2] PENSACK, D. A. & COREY, E. J. Lhasa—logic and heuristics applied to synthetic analysis. In *Computer-Assisted Organic Synthesis*, vol. 61 of *ACS Symposium Series*, 1–32 (AMERICAN CHEMICAL SOCIETY, 1977). URL <https://doi.org/10.1021/bk-1977-0061.ch001>.
- [3] Avila, C. *et al.* Chemistry in a graph: modern insights into commercial organic synthesis planning. *Digital Discovery* **3**, 1682–1694 (2024). URL <http://dx.doi.org/10.1039/D4DD00120F>.
- [4] Levin, I., Liu, M., Voigt, C. A. & Coley, C. W. Merging enzymatic and synthetic chemistry with computational synthesis planning. *Nature Communications* **13**, 7747 (2022).
- [5] Nicolaou, K. & Montagnon, T. Molecules that changed the world. *Clinical Pharmacology and Pharmacokinetics* **14** (2008).
- [6] Brown, D. G. & Boström, J. Analysis of past and present synthetic methodologies on medicinal chemistry: Where have all the new reactions gone? *Journal of Medicinal Chemistry* **59**, 4443–4458 (2016).
- [7] Hart, P. E., Nilsson, N. J. & Raphael, B. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* **4**, 100–107 (1968).
- [8] Schwaller, P. *et al.* Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *Chemical Science* **10**, 369–377 (2019).
- [9] Kreutter, D., Schwaller, P. & Reymond, J.-L. Predicting enzymatic reactions with a molecular transformer. *Chem. Sci.* **12**, 8648–8659 (2021). URL <http://dx.doi.org/10.1039/D1SC02362D>.
- [10] Tetko, I. V., Karpov, P., Deursen, R. V. & Godin, G. State-of-the-art augmented nlp transformer models for direct and single-step retrosynthesis. *Nature Communications* **11**, 5575 (2020). URL <https://doi.org/10.1038/s41467-020-19266-y>.
- [11] Irwin, R., Dimitriadis, S., He, J. & Bjerrum, E. J. Chemformer: a pre-trained transformer for computational chemistry. *Machine Learning: Science and Technology* **3**, 015022 (2022). URL <https://dx.doi.org/10.1088/2632-2153/ac3ffb>.
- [12] Granqvist, E., Mercado, R. & Genheden, S. Retrosynformer: planning multi-step chemical synthesis routes via a decision transformer. *Digital Discovery* – (2025). URL <http://dx.doi.org/10.1039/D5DD00153F>.
- [13] Wang, X. *et al.* Retroprime: A diverse, plausible and transformer-based method for single-step retrosynthesis predictions. *Chemical Engineering Journal* **420**, 129845 (2021). URL <https://www.sciencedirect.com/science/article/pii/S1385894721014303>.
- [14] Andronov, M., Andronova, N., Wand, M., Schmidhuber, J. & Clevert, D.-A. Fast and scalable retrosynthetic planning with a transformer neural network and speculative beam search. *arXiv* **abs/2508.01459** (2025).
- [15] Bradshaw, J., Kusner, M. J., Paige, B., Segler, M. H. S. & Hernández-Lobato, J. M. A generative model for electron paths. In *International Conference on Learning Representations* (2019). URL <https://openreview.net/forum?id=r1x4BnCcKX>.
- [16] Sacha, M. *et al.* Molecule edit graph attention network: Modeling chemical reactions as sequences of graph edits. *Journal of Chemical Information and Modeling* **61**, 3273–3284 (2021). URL <https://doi.org/10.1021/acs.jcim.1c00537>.
- [17] Coley, C. *et al.* A graph-convolutional neural network model for the prediction of chemical reactivity. *Chem. Sci.* **10**, 370–377 (2019). URL <http://dx.doi.org/10.1039/C8SC04228D>.
- [18] Qian, W. W. *et al.* Integrating deep neural networks and symbolic inference for organic reactivity prediction (2020).
- [19] Bi, H. *et al.* Non-autoregressive electron redistribution modeling for reaction prediction. In Meila, M. & Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning*, vol. 139 of *Proceedings of Machine Learning Research*, 904–913 (PMLR, 2021). URL <https://proceedings.mlr.press/v139/bi21a.html>.
- [20] Corey, E., Long, A. & Rubenstein, S. Computer-assisted analysis in organic synthesis. *Science* **228**, 408–418 (1985).
- [21] Segler, M. H. S. & Waller, M. P. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry – A European Journal* **23**, 5966–5971 (2017). URL <https://chemistry-europe.onlinelibrary.wiley.com/doi/abs/10.1002/chem.201605499>. <https://chemistry-europe.onlinelibrary.wiley.com/doi/pdf/10.1002/chem.201605499>.



- [22] Coley, C. W., Barzilay, R., Jaakkola, T. S., Green, W. H. & Jensen, K. F. Prediction of organic reaction outcomes using machine learning. *ACS Central Science* **3**, 434–443 (2017). URL <https://doi.org/10.1021/acscentsci.7b00064>.
- [23] Zhang, X., Lin, H., Zhang, M., Zhou, Y. & Ma, J. A data-driven group retrosynthesis planning model inspired by neurosymbolic programming. *Nature Communications* **16** (2025).
- [24] Roh, J. *et al.* Higher-level strategies for computer-aided retrosynthesis. *ChemRxiv* (2025).
- [25] Chen, S. & Jung, Y. A generalized-template-based graph neural network for accurate organic reactivity prediction. *Nature Machine Intelligence* **4**, 772–780 (2022). URL <https://doi.org/10.1038/s42256-022-00526-z>.
- [26] Seidl, P. *et al.* Improving few-and zero-shot reaction template prediction using modern hopfield networks. *Journal of chemical information and modeling* **62**, 2111–2120 (2022).
- [27] Choe, J., Kim, H., Chok, Y. T., Gim, M. & Kang, J. Retrosynthetic crosstalk between single-step reaction and multi-step planning. *Journal of Cheminformatics* **17** (2025).
- [28] Wang, X. *et al.* Towards efficient discovery of green synthetic pathways with monte carlo tree search and reinforcement learning. *Chem. Sci.* **11**, 10959–10972 (2020). URL <http://dx.doi.org/10.1039/D0SC04184J>.
- [29] Chen, B., Li, C., Dai, H. & Song, L. Retro*: Learning retrosynthetic planning with neural guided a* search. In *The 37th International Conference on Machine Learning (ICML 2020)* (2020).
- [30] Yu, K. *et al.* Double-ended synthesis planning with goal-constrained bidirectional search (2024). URL <https://arxiv.org/abs/2407.06334>. 2407.06334.
- [31] Finnigan, W., Hepworth, L. J., Flitsch, S. L. & Turner, N. J. Retrobiocat as a computer-aided synthesis planning tool for biocatalytic reactions and cascades. *Nature catalysis* **4**, 98–104 (2021).
- [32] Sheldon, R. A. The e factor: fifteen years on. *Green Chem.* **9**, 1273–1283 (2007). URL <http://dx.doi.org/10.1039/B713736M>.
- [33] Horváth, I. T. & Anastas, P. T. Innovations and green chemistry. *Chemical Reviews* **107**, 2169–2173 (2007). URL <https://doi.org/10.1021/cr078380v>.
- [34] Ramsauer, H. *et al.* Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217* (2020).
- [35] Genheden, S. & Bjerrum, E. Paroutes: towards a framework for benchmarking retrosynthesis route predictions. *Digital Discovery* **1**, 527–539 (2022). URL <http://dx.doi.org/10.1039/D2DD00015F>.
- [36] Draghici, C. & Njardarson, J. T. Chemistry by design: A web-based educational flashcard for exploring synthetic organic chemistry. *Journal of Chemical Education* **89**, 1080–1082 (2012). URL <https://doi.org/10.1021/ed2006423>.
- [37] Lang, M., Stelzer, M. & Schomburg, D. Bkm-react, an integrated biochemical reaction database. *BMC biochemistry* **12**, 1–9 (2011).
- [38] Bansal, P. *et al.* Rhea, the reaction knowledgebase in 2022. *Nucleic Acids Research* **50**, D693–D700 (2021). URL <https://doi.org/10.1093/nar/gkab1016>. <https://academic.oup.com/nar/article-pdf/50/D1/D693/42058388/gkab1016.pdf>.
- [39] Landrum, G. Rdkit: Open-source cheminformatics software (2016). URL https://github.com/rdkit/rdkit/releases/tag/Release_2016_09_4.
- [40] Lowe, D. Chemical reactions from US patents (1976-Sep2016) (2017). URL https://figshare.com/articles/dataset/Chemical_reactions_from_US_patents_1976-Sep2016_/5104873.
- [41] Mary, H. *et al.* datamol-io/datamol: 0.12.3 (2024). URL <https://doi.org/10.5281/zenodo.10535844>.
- [42] Glorot, X. & Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In Teh, Y. W. & Titterton, M. (eds.) *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, vol. 9 of *Proceedings of Machine Learning Research*, 249–256 (PMLR, Chia Laguna Resort, Sardinia, Italy, 2010). URL <https://proceedings.mlr.press/v9/glorot10a.html>.
- [43] Ioffe, S. & Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456 (pmlr, 2015).
- [44] Paszke, A. *et al.* *PyTorch: an imperative style, high-performance deep learning library* (Curran Associates Inc., Red Hook, NY, USA, 2019).
- [45] Coley, C. W. *et al.* A robotic platform for flow synthesis of organic compounds informed by ai planning. *Science* **365**, eaax1566 (2019).
- [46] Molport. Easy compound ordering service - molport. <https://www.molport.com/shop/index>.



- [47] Kiss, R., Sandor, M. & Szalai, F. A. <http://mcule.com>: a public web service for drug discovery. *Journal of Cheminformatics* **4**, P17 (2012). URL <https://doi.org/10.1186/1758-2946-4-S1-P17>. Copyright 2012 Kiss et al; licensee BioMed Central Ltd.
- [48] ChemSpace. ChemSpace. <https://chem-space.com>. Accessed: 2024-02-21.
- [49] Sorkun, M. C., Saliou, B. & Er, S. Chemprice, a python package for automated chemical price search. ChemRxiv (2024).
- [50] Sterling, T. & Irwin, J. J. Zinc 15 – ligand discovery for everyone. *Journal of Chemical Information and Modeling* **55**, 2324–2337 (2015). URL <https://doi.org/10.1021/acs.jcim.5b00559>.
- [51] Gao, H. *et al.* Using machine learning to predict suitable conditions for organic reactions. *ACS central science* **4**, 1465–1476 (2018).
- [52] Razzaq, T. & Kappe, C. Continuous flow organic synthesis under high-temperature/pressure conditions. *Chemistry – An Asian Journal* **5**, 1274–1289 (2010). URL <https://aces.onlinelibrary.wiley.com/doi/abs/10.1002/asia.201000010>. <https://aces.onlinelibrary.wiley.com/doi/pdf/10.1002/asia.201000010>.
- [53] ACS Green Chemistry Institute® Pharmaceutical Roundtable. Solvent selection guide: Version 2.0 (2011). Retrieved 12th May 2024 from: <http://www.acs.org/content/acs/en/greenchemistry/research-innovation/tools-for-green-chemistry.html>.
- [54] Gallo, K. *et al.* SuperNatural 3.0—a database of natural products and natural product-based derivatives. *Nucleic Acids Research* **51**, D654–D659 (2022). URL <https://doi.org/10.1093/nar/gkac1008>. <https://academic.oup.com/nar/article-pdf/51/D1/D654/48440479/gkac1008.pdf>.
- [55] Wishart, D. *et al.* T3DB: the toxic exposome database. *Nucleic Acids Research* **43**, D928–D934 (2014). URL <https://doi.org/10.1093/nar/gku1004>. <https://academic.oup.com/nar/article-pdf/43/D1/D928/7311219/gku1004.pdf>.
- [56] Wang, J. *et al.* Enantioselective synthesis of the 1, 3-dienyl-5-alkyl-6-oxy motif: Method development and total synthesis. *Angewandte Chemie International Edition* **63**, e202400478 (2024).
- [57] Genheden, S. *et al.* Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of Cheminformatics* **12**, 70 (2020). URL <https://doi.org/10.1186/s13321-020-00472-1>.
- [58] Yang, Y. *et al.* Batgpt-chem: A foundation large model for chemical engineering. *Research* **8**, 0827 (2025). URL <https://spj.science.org/doi/abs/10.34133/research.0827>. <https://spj.science.org/doi/pdf/10.34133/research.0827>.
- [59] Zhao, Z. *et al.* Developing chemdfm as a large language foundation model for chemistry. *Cell Reports Physical Science* **6**, 102523 (2025). URL <https://www.sciencedirect.com/science/article/pii/S2666386425001225>.



The synthetic dataset utilized in this study was obtained from the USPTO40 repository, while the enzymatic dataset was sourced from the BKMS37 and RHEA38 databases. All datasets are publicly accessible and open-source. The processed datasets generated and analyzed during the study and the trained model weights are available on Figshare.

