

Self-Improved Retrosynthetic Planning

Junsu Kim¹ Sungsoo Ahn² Hankook Lee¹ Jinwoo Shin¹

Abstract

Retrosynthetic planning is a fundamental problem in chemistry for finding a pathway of reactions to synthesize a target molecule. Recently, search algorithms have shown promising results for solving this problem by using deep neural networks (DNNs) to expand their candidate solutions, i.e., adding new reactions to reaction pathways. However, the existing works on this line are suboptimal; the retrosynthetic planning problem requires the reaction pathways to be (a) represented by real-world reactions and (b) executable using “building block” molecules, yet the DNNs expand reaction pathways without fully incorporating such requirements. Motivated by this, we propose an end-to-end framework for directly training the DNNs towards generating reaction pathways with the desirable properties. Our main idea is based on a self-improving procedure that trains the model to imitate successful trajectories found by itself. We also propose a novel reaction augmentation scheme based on a forward reaction model. Our experiments demonstrate that our scheme significantly improves the success rate of solving the retrosynthetic problem from 86.84% to 96.32% while maintaining the performance of DNN for predicting valid reactions.

1. Introduction

To synthesize a novel molecule, chemists require executing a pathway of reactions starting from a set of known or commercially available building block molecules. Hence, discovering such a reaction pathway for a target molecule is crucial in important applications such as drug discovery (Hughes et al., 2011) and material design (Yan et al., 2018).

To tackle this problem, retrosynthetic planning (Corey,

¹Korea Advanced Institute of Science and Technology (KAIST) ²Mohamed bin Zayed University of Artificial Intelligence (MBZUAI). Correspondence to: Junsu Kim <junsu.kim@kaist.ac.kr>.

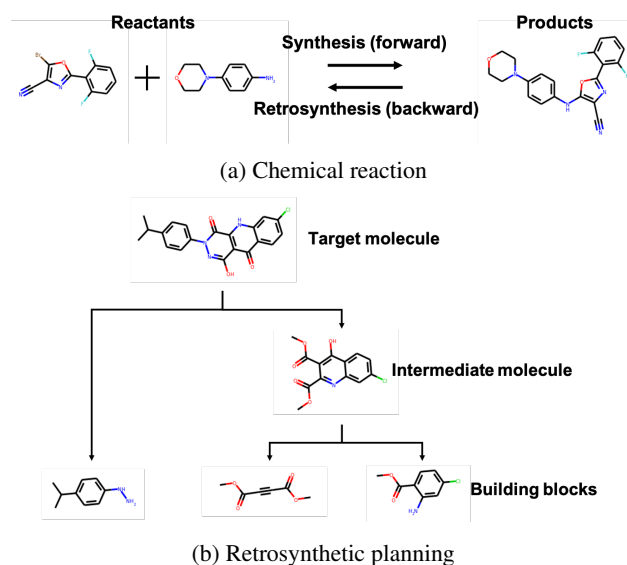


Figure 1. Illustration of (a) synthesis (forward) and retrosynthesis (backward) with respect to a chemical reaction and (b) outcome of retrosynthetic planning given the target molecule 2-(2,6-difluorophenyl)-5-(4-morpholin-4-ylanilino)-1,3-oxazole-4-carbonitrile. Given a target molecule, a retrosynthetic planning algorithm aims at finding a reaction pathway ending up in the building block molecules.

1991) finds a series of chemically valid reactions starting from the target molecule until reaching the building block molecules in a backward and recursive manner. The main challenge of retrosynthetic planning is twofold: (a) finding an accurate single-step retrosynthetic model that predicts a single reaction of a given product and (b) designing an efficient search algorithm for a reaction pathway starting from the set of building block molecule.

Especially, recent works have proposed deep neural networks (DNNs) as attractive models for single-step retrosynthesis. Using the existing real-world reaction datasets (Lowe, 2012), they train (in a supervised manner) and evaluate DNNs to predict a reactant-set from a given product. To be specific, existing works use DNNs to predict the reactant-set based on applying a known reaction template to the product (Segler & Waller, 2017; Dai et al., 2019), generating each reactant from scratch (Liu et al., 2017; Karpov

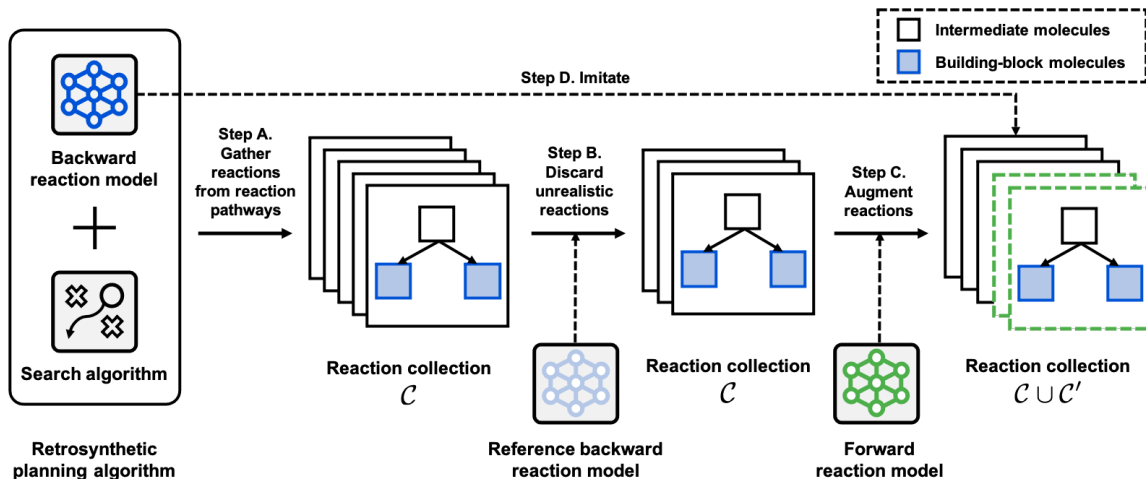


Figure 2. Illustration of our framework. Our framework iterates the four-step procedure. In step A, we gather reactions from the reaction pathways that are generated via a search algorithm combined with a backward reaction model to form a collection of reactions \mathcal{C} . In step B, we discard unrealistic reactions in the collection \mathcal{C} using a reference backward model. In step C, we generate a set of reactions \mathcal{C}' from augmenting the reactions in the collection \mathcal{C} using a forward reaction model. In step D, we train the backward reaction model to imitate reactions in the collection $\mathcal{C} \cup \mathcal{C}'$.

et al., 2019; Zheng et al., 2019), or modifying the product using atom-wise and bond-wise operations (Shi et al., 2020; Somnath et al., 2020; Yan et al., 2020).

On the other hand, researchers also have developed efficient search algorithms for retrosynthetic planning based on the DNN-based single-step retrosynthetic models. Their main idea is to represent retrosynthetic planning as a sequential decision making problem and apply tree search algorithms such as Monte Carlo tree search (Segler et al., 2018), proof number search (Kishimoto et al., 2019), and A* search (Chen et al., 2020b).

Intriguingly, most of the existing DNN-based retrosynthetic planning frameworks are not end-to-end. Namely, the performance of a retrosynthetic planning algorithm can be evaluated by (a) whether if the algorithm proposes reaction pathways representing reactions existing in real-world and (b) the success rate of finding such reaction pathways starting from the set of building block molecules. Since existing frameworks optimize DNN-based single-step retrosynthetic models and search algorithms for (a) and (b) separately, they may have suboptimal performance.

Contribution. In this paper, we propose a new end-to-end framework for retrosynthetic planning based on training the DNN-based single-step retrosynthetic model toward maximizing the performance of retrosynthetic planning. We train DNNs for maximizing the success rate of search algorithms in addition to representing the inverse of real-world reactions. While our framework can be simply implemented on top of existing frameworks for retrosynthetic planning, we empirically observe that our end-to-end training of DNN

leads to surprisingly large performance gains.

To train the single-step retrosynthetic model for maximizing the success rate of search algorithms, we introduce a self-improving procedure that trains the model to imitate successful trajectories found by itself combined with the search algorithm. To train the model to generate realistic reaction pathways, we additionally introduce a likelihood-based criterion for filtering out samples used in the self-improving algorithm. Finally, to improve the generalization ability of the single-step retrosynthetic model, we propose a novel augmentation scheme based on modifying reactions using a forward reaction model. We provide an overall illustration of our framework in Figure 2.

To demonstrate the effectiveness of our framework, we conduct experiments based on the USPTO database (Lowe, 2012). Thanks to imitating reactions that are realistic and executable from building block molecules, our framework significantly improves the success rate of solving the retrosynthetic problem from 86.84% to 96.32%. Moreover, the reduced average time for planning demonstrates the efficacy of our framework. The average length and cost of searched pathways, which is a metric to measure the quality, also decreased than other baselines. In our ablation studies, we show the effectiveness of each component in our framework.

Our work reduces the gap between the widely used supervised learning of single-step retrosynthetic models and the goal of retrosynthetic planning. We believe that our work would guide new interesting directions in the future by bridging this gap.

2. Preliminary

2.1. Problem Setup

The task of retrosynthetic planning is to search for a set of chemical reactions $\tau = \{R_i\}_{i=1}^N$, i.e., a reaction pathway, required for synthesizing a target molecule t . Each reaction $R = (m, \mathcal{R})$ is represented by a pair of a product m and a reactant-set $\mathcal{R} = \{r_j\}_{j=1}^M$.¹ Furthermore, the reaction pathways are desired to satisfy the following conditions:

- A. The reactions should correspond to a realistic pair of a product and a reactant-set, i.e., any reaction R should be executable in the real world.
- B. Any reactant r in a reaction $R \in \tau$ should be either a member of building block molecules \mathcal{I} or a product of another reaction $R' \in \tau$.

For training and evaluating retrosynthetic planning algorithms, we assume having access to datasets $\mathcal{D}_{\text{target}}$ and $\mathcal{D}_{\text{reaction}}$ consisting of real-world target molecules and reactions, respectively.

2.2. Forward and Backward Reaction Models

We consider retrosynthetic planning based on deep neural networks (DNNs) for representing a *backward reaction model* $p_b(\mathcal{R}|m; \theta_b)$ with parameter θ_b , i.e., a single-step retrosynthetic model. In addition, we consider a *forward reaction model* $p_f(m|\mathcal{R}; \theta_f)$ with parameter θ_f . Such a choice allows the reaction models to flexibly incorporate the complex chemical knowledge based on the expressive power of DNNs.

Researchers have developed various ways of modeling forward reaction (Jin et al., 2017; Bradshaw et al., 2018; Schwaller et al., 2018; Do et al., 2019; Schwaller et al., 2019), and backward reaction (Segler & Waller, 2017; Dai et al., 2019; Liu et al., 2017; Karpov et al., 2019; Zheng et al., 2019; Shi et al., 2020; Somnath et al., 2020; Yan et al., 2020) using DNNs. They are mainly categorized into either template-based or template-free approach depending on their reliance on the reaction templates, i.e., subgraph patterns describing how the chemical reaction occurs among reactants.

In this work, we consider template-based reaction models represented by a multi-layered perceptron (MLP), prioritizing templates to apply in the given templates list (Segler & Waller, 2017). To be specific, the MLP is trained to predict plausible templates to apply, taking Morgan fingerprint (Rogers & Hahn, 2010), a fixed-size vector representation of a set of molecules, as an input.

¹To simplify the problem, we omit other conditions for describing a reaction, e.g., reagents. However, they can be incorporated into our framework with relatively small modifications.

2.3. Search Algorithms

To solve retrosynthetic planning with respect to the exponentially large space of reaction pathways, it is crucial to use an efficient *search algorithm* \mathcal{A} . Existing search algorithms (Segler et al., 2018; Kishimoto et al., 2019; Chen et al., 2020b) build reaction pathways by updating them in a backward direction, i.e., adding reactions for synthesizing a product in the intermediate reaction pathway. They typically use a backward reaction model trained on a real-world dataset to propose a reaction pathway consisting of realistic reactions. The data structure and the protocol for updating the intermediate reaction pathways are specific to search algorithms, e.g., Monte Carlo tree search (Segler & Waller, 2017) and proof number search (Kishimoto et al., 2019).

To be specific, Segler et al. (2018) design a search tree where each node represents a set of molecules and expand the tree by balancing the selection of high-value nodes and unexplored nodes. The value of each node is estimated based on a Monte Carlo rollout of backward reaction models. Furthermore, Kishimoto et al. (2019) and Chen et al. (2020b) employ AND-OR search trees to represent retrosynthetic pathways, where OR and AND nodes represent molecules and reactions, respectively. The most promising node to expand during planning is selected using human designed heuristics (Kishimoto et al., 2019) or a DNN trained on an offline dataset (Chen et al., 2020b).

In this work, we consider the recently proposed RETRO* (Chen et al., 2020b) for traversing the space of reaction pathways since it has demonstrated strong performance. The RETRO* algorithm mimics the A* algorithm by performing the best-first search based on the cost of the current path and the estimated cost to the goal. They consider the cost of the current path as the sum of reaction costs, and the estimated cost to the goal is computed from the value function, which is parameterized by a neural network trained using an existing dataset of reactions. Additionally, Chen et al. (2020b) also introduce RETRO*-0, which does not utilize the value function for expansion that makes a tradeoff between its performance and the expense of using an additional DNN for representing the value function.

3. Self-Improved Retrosynthetic Planning

3.1. Overview of Self-Improved Retrosynthetic Planning

In this section, we introduce a new framework for retrosynthetic planning based on a self-improved model adaptation procedure. Similar to prior works, our framework aims to find reaction pathways by running a search algorithm \mathcal{A} using reactions suggested by a backward reaction model $p_b(m|\mathcal{R}; \theta_b)$ with parameter θ_b . However, our framework differs from the existing works by adapting the backward re-

action model towards improving the performance of search algorithms; existing works use backward reaction models that are agnostic to the choice of search algorithm and may yield suboptimal performance.

At a high level, our framework trains the backward reaction model $p_b(m|\mathcal{R};\theta_b)$ to maximize the likelihood of generated reaction pathways from itself combined with search algorithm \mathcal{A} . To improve the quality of reactions from the reaction pathways used for imitation learning, we introduce a *reference backward reaction model* $p_b(m|\mathcal{R};\bar{\theta}_b)$ using the same architecture as the original backward reaction model. It is trained on a real-world reaction dataset $\mathcal{D}_{\text{reaction}}$ to let its likelihood determine whether a reaction resembles reactions existing in the real world. Furthermore, we propose a novel reaction augmentation scheme based on a forward reaction model $p_f(m|\mathcal{R};\theta_f)$ to improve the diversity of reactions used for imitation learning. To further clarify our objective, note that we aim to optimize both the success rate of planning and the synthesis route cost. Following the prior work (Chen et al., 2020b), we consider $-\sum_{(m,\mathcal{R})\in\tau}\log p_b(\mathcal{R}|m;\bar{\theta}_b)$ as the synthesis route cost: accumulation of negative log-likelihoods of reactions under the reference backward model.

Our framework repeats the following steps:

- **Step A.** Generate a set of reaction pathways based on combining a search algorithm with the backward reaction model. Gather reactions from the reaction pathways to form a collection of reactions \mathcal{C} .
- **Step B.** Discard any reactions from the collection \mathcal{C} that are determined to be unrealistic using a reference backward reaction model.
- **Step C.** Generate a set of reactions \mathcal{C}' from augmenting the reactions in the collection \mathcal{C} by replacing the corresponding product using a forward reaction model.
- **Step D.** Train the backward reaction model by maximizing the log-likelihood of the reactions in $\mathcal{C}\cup\mathcal{C}'$, i.e., maximize $\sum_{(m,\mathcal{R})\in\mathcal{C}\cup\mathcal{C}'}\log p_b(\mathcal{R}|m;\theta_b)$.

To speed up training, we initialize the backward reaction model $p_b(m|\mathcal{R};\theta_b)$ with supervised learning on the reaction dataset $\mathcal{D}_{\text{reaction}}$. Intuitively, our algorithm allows training the backward reaction model on reactions with their quality improved by using the search algorithm, i.e., RETRO*. Since we additionally filter out unrealistic reactions based on the reference backward reaction model, our backward reaction model retains the ability to generate realistic reactions.

Our self-improving procedure is similar to the prior works such as DAgger (Ross et al., 2011), retrospective imitation learning (Song et al., 2018), expert iteration (Anthony et al.,

Algorithm 1 Self-Improved Retrosynthetic Planning

Input: backward reaction model p_b , forward reaction model p_f , retrosynthetic planning algorithm \mathcal{A} , target molecule dataset $\mathcal{D}_{\text{target}}$, reaction dataset $\mathcal{D}_{\text{reaction}}$, and filtering thresholds $\epsilon, \epsilon_{\text{aug}}$.

Maximize $\sum_{(m,\mathcal{R})\in\mathcal{D}_{\text{reaction}}} p_b(\mathcal{R}|m;\theta_b)$ over θ_b .

Maximize $\sum_{(m,\mathcal{R})\in\mathcal{D}_{\text{reaction}}} p_b(\mathcal{R}|m;\bar{\theta}_b)$ over $\bar{\theta}_b$.

Maximize $\sum_{(m,\mathcal{R})\in\mathcal{D}_{\text{reaction}}} p_f(m|\mathcal{R};\theta_f)$ over θ_f .

for $i = 1, \dots, I$ **do**

 Initialize a collection of reaction pathways $\mathcal{C} \leftarrow \emptyset$.

for $n = 1, \dots, N$ **do**

 Sample a molecule $t \sim \mathcal{D}_{\text{target}}$.

 Compute a reaction pathway $\tau \leftarrow \mathcal{A}(t, p_b)$.

for $(m, \mathcal{R}) \in \tau$ **do**

if $p_b(\mathcal{R}|m;\bar{\theta}_b) > \epsilon$ **then**

 Update the collection $\mathcal{C} \leftarrow \mathcal{C} \cup \{(m, \mathcal{R})\}$.

end if

end for

end for

 Initialize a collection of reaction pathways $\mathcal{C}' \leftarrow \emptyset$.

for $(m, \mathcal{R}) \in \mathcal{C}$ **do**

 Compute $m' \leftarrow \arg \max_m p_f(m|\mathcal{R};\theta_f)$.

if $p_f(m'|\mathcal{R};\theta_f) > \epsilon_{\text{aug}}$ **and**

$\mathcal{R} = \arg \max_{\mathcal{R}} p_b(\mathcal{R}|m';\bar{\theta}_b)$ **then**

 Update the collection $\mathcal{C}' \leftarrow \mathcal{C}' \cup \{(m', \mathcal{R})\}$.

end if

end for

 Maximize $\sum_{(m,\mathcal{R})\in\mathcal{C}\cup\mathcal{C}'} \log p_b(\mathcal{R}|m;\theta_b)$ over θ_b .

end for

2017; 2019), AlphaGo Zero (Silver et al., 2017), and NEXT (Chen et al., 2020a) in meta path planning. Meanwhile, one can note that RetroGNN (Liu et al., 2020) uses routes found by a retrosynthetic planning algorithm to learn a value function, whereas our framework learns a policy (backward reaction model).

We also note that Schreck et al. (2019) proposed an end-to-end framework based on reinforcement learning of the backward reaction model towards maximizing the success rate of finding a reaction pathway for the target molecule. However, they do not consider whether if the models propose realistic reaction pathways and are not directly comparable with our work.

We provide an illustration and a detailed description of our framework in Figure 2 and Algorithm 1, respectively.

3.2. Detailed Components of Self-Improved Retrosynthetic Planning

In the rest of this section, we provide a detailed description of our algorithmic components: generating reaction pathways using a search algorithm, evaluating realistic-ness of




Figure 3. Example of (a) extracted reaction from reaction pathways, and (b) corresponding augmented reaction via the forward reaction model.

reactions using a reference backward reaction model, and augmenting reactions using a forward reaction model.

Generating reaction pathways. To generate reaction pathways, we sample a target molecule t from the target molecule dataset $\mathcal{D}_{\text{target}}$ and apply the search algorithm \mathcal{A} based on the current backward reaction model $p_b(\mathcal{R}|m; \theta_b)$. In this work, we consider RETRO* (Chen et al., 2020b) for generating the reaction pathways. However, our framework is general and applicable to other existing search algorithms such as Monte Carlo tree search (Segler et al., 2018) and proof number search (Kishimoto et al., 2019).

Filtering out unrealistic reactions. To prevent the backward reaction model from learning to predict unrealistic reactions, we use a reference backward reaction model $p_b(\mathcal{R}|m; \theta_b)$ to determine whether a reaction is realistic or not. To be specific, we train the reference backward model on a real-world reaction dataset $\mathcal{D}_{\text{reaction}}$ and use its likelihood to measure the realistic-ness of the reactions; we filter out generated reactions whose likelihood under the reference backward model is less than ϵ .

Reaction augmentation. The generalization capability of the backward reaction model can be improved by training on a diverse set of reactions. To this end, we propose a new augmentation scheme based on domain knowledge of chemistry; there can be multiple products resulting from the same reactant-set. Based on this knowledge, we augment the existing reaction $R = (m, \mathcal{R})$ by replacing the product m with a new product m' proposed by a forward reaction model $p_f(\cdot|\mathcal{R}; \theta_f)$. In order to augment the reactions in a realistic way, we additionally filter out the proposed product m' when it is not confident from the forward reaction model or the reference backward reaction model, i.e., we reject m' when $p_f(m'|\mathcal{R}; \theta_f) \leq \epsilon_{\text{aug}}$ or $\mathcal{R} \neq \arg \max_{\mathcal{R}} p_b(\mathcal{R}|m'; \theta_b)$. We note that our augmentation method is computationally cheaper than generating

reactions from scratch. We demonstrate an example of outcomes from our augmentation scheme in Figure 3.

4. Experiments

4.1. Experimental Setup

Our framework requires specifying (1) a set of building block molecules \mathcal{I} , (2) a target molecule dataset $\mathcal{D}_{\text{target}}$, (3) a reaction dataset $\mathcal{D}_{\text{reaction}}$, (4) a retrosynthetic planning algorithm \mathcal{A} , and (5) a backward reaction model p_b , a reference backward reaction model \bar{p}_b , a forward reaction model p_f . The details of the components are as follows.

Dataset. For building blocks \mathcal{I} , we use all of 231M commercially available molecules present in *eMolecules*.² Note that chemists can choose another set of molecules for \mathcal{I} depending on their own circumstances, such as a financial budget. For the target molecules $\mathcal{D}_{\text{target}}$, we choose synthesizable molecules from \mathcal{I} and reactions in the United States Patent Office (USPTO) database (Lowe, 2012). To this end, we follow the procedure described by Chen et al. (2020b) and then obtain 299202 target molecules. For the reaction dataset $\mathcal{D}_{\text{reaction}}$, we use reactions extracted from USPTO, following training/validation/test splits by Chen et al. (2020b).

Retrosynthetic planning algorithm. We use RETRO* and RETRO*-0 (Chen et al., 2020b) as the search algorithm \mathcal{A} in our framework. RETRO*-0 denotes its variant not relying on the value function, which estimates the cost to synthesize a given molecule, proposed by Chen et al. (2020b).

Model. We use template-based 2-layer MLPs (Segler & Waller, 2017) for parameterizing backward reaction model p_b , reference backward reaction model \bar{p}_b , and forward reaction model p_f . More specifically, the models predict a plausible reaction template among a pre-defined set of templates \mathcal{T} , i.e., they can be considered as classification models. We use RDChiral (Coley et al., 2019) for extracting reaction templates from the USPTO database, and it results in 380k reaction templates. We use the Morgan fingerprint (Rogers & Hahn, 2010), of radius 2 with 2048 bits, as an input of MLPs.

Training. The reference backward reaction model \bar{p}_b and forward reaction model p_f are trained using reactions in the dataset $\mathcal{D}_{\text{reaction}}$ and then frozen before conducting our framework. Instead of training a reference backward reaction model from scratch, we use the backward reaction model trained by Chen et al. (2020b)³ for our reference backward model \bar{p}_b . The forward reaction model p_f is

²<http://downloads.emolecules.com/free/2019-11-01/>

³https://github.com/binghong-ml/retro_star

Table 1. Performance of backward reaction model and retrosynthetic planning in the USPTO dataset. The backward reaction model is evaluated using TOP-1 and TOP-10 exact match accuracy (%). Retrosynthetic planning is evaluated using SUCC. RATE ($N = 50$), SUCC. RATE ($N = 500$), LENGTH, TIME, and COST. N denotes the limit of backward reaction model calls. LENGTH is the number of reactions in a route. TIME is measured by the number of backward reaction model calls, with a hard limit of 500. †The experimental results of GREEDY DFS, MCTS, and DFPN-E are from Chen et al. (2020b). The best results are marked in bold. We use brackets to report the relative gains over each counterpart that does not use our framework.

ALGORITHM	REACTIONS		REACTION PATHWAYS				
	TOP-1 \uparrow	TOP-10 \uparrow	SUCC. RATE \uparrow ($N = 50$)	SUCC. RATE \uparrow ($N = 500$)	LENGTH \downarrow	TIME \downarrow	COST \downarrow
GREEDY DFS [†]	-	-	-	22.63	-	388.15	-
MCTS [†]	-	-	-	33.68	-	370.51	-
DFPN-E [†]	-	-	-	55.26	-	279.67	-
RETRO*-0	44.53	72.71	27.37	79.47	11.21	208.09	19.40
RETRO*-0 + OURS	44.03 (-1.12%)	73.14 (+0.59%)	57.37 (+109.62%)	96.32 (+21.20%)	7.69 (-31.40%)	96.22 (-53.76%)	11.66 (-39.90%)
RETRO*	44.53	72.71	44.21	86.84	9.71	157.11	15.33
RETRO* + OURS	44.03 (-1.12%)	73.15 (+0.61%)	57.89 (+30.94%)	91.05 (+4.85%)	8.74 (-9.99%)	100.15 (-36.25%)	15.23 (-0.65%)

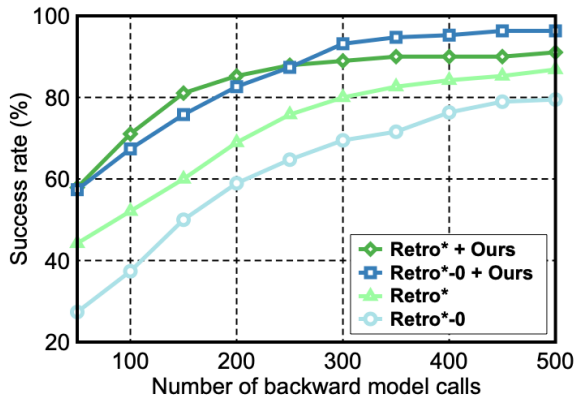


Figure 4. Success rate (%) under varying limits of backward reaction model calls. Our framework outperforms the best baselines, RETRO*-0, regardless of the limits.

trained with a learning rate of 0.001 for 100 epochs. The parameters of the backward reaction model p_b are initialized to that of the reference backward reaction model \bar{p}_b . During the self-improving procedure in our framework, the backward reaction model p_b is trained with a learning rate of 0.0001 for 20 epochs. Adam optimizer (Kingma & Ba, 2014) is used with a mini-batch of size 1024 for training all the models. We iterate our overall procedure three times.

Filtering thresholds. In our framework, there exists two thresholds: (1) ϵ for removing unrealistic reactions from success routes via a reference backward reaction model, (2) ϵ_{aug} for rejecting unconfident augmented reaction generated

Table 2. Ablation study on augmentation via the forward reaction model. SEARCH indicates the type of search algorithm \mathcal{A} . AUG indicates the augmentation via the forward reaction model. N denotes the limit of backward reaction model calls.

SEARCH	AUG	SUCC. RATE \uparrow		
		$N = 50$	$N = 250$	$N = 500$
RETRO*-0		46.84	81.05	92.63
RETRO*-0	✓	50.00	83.16	93.16
RETRO*		51.58	82.11	88.95
RETRO*	✓	51.58	83.68	90.00

by a forward reaction model. We set both thresholds ϵ , ϵ_{aug} as 0.8. Namely, we filter out reactions of which likelihood under the corresponding model is under 0.8.

Evaluation. We evaluate retrosynthetic planning for 190 target molecules in a limited time budget (i.e., the number of calls of backward reaction model p_b) following Chen et al. (2020b). In the limit, we measure the success rate, the average time of planning, the average length, and the cost of discovered routes. Note that we consider the cost of a route as the summation of negative log-likelihoods of reactions in the route following Chen et al. (2020b), i.e., $-\sum_{(m, \mathcal{R}) \in \tau} \log p_b(\mathcal{R}|m; \bar{\theta}_b)$. We evaluate the backward reaction model using widely-used top- k exact match accuracy in the test split of the reaction dataset $\mathcal{D}_{\text{reaction}}$.

Baselines. To demonstrate the effectiveness of our framework, we compare our method with existing retrosynthetic planning frameworks such as GREEDY DFS, MCTS (Segler

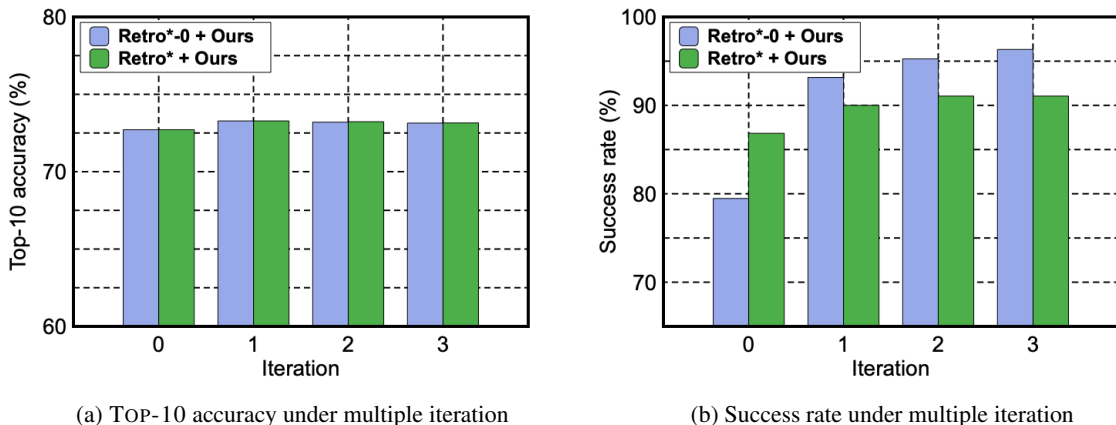


Figure 5. We repeat our procedure multiple times and investigate the performance of the backward reaction model and the retrosynthetic planning. Iteration 0 is vanilla RETRO*-0 or RETRO*, which our framework is not applied yet. As we iterate our framework, we can further improve the performance of retrosynthetic planning while maintaining the reliability of the backward reaction model.

Table 3. Experimental results of our framework with different filtering thresholds ϵ in reaction extraction step using the reference backward model. As we do not filter out any reactions from success routes, we suffer performance degradation in TOP-1 and TOP-10 accuracy, as unrealistic reactions can be included in the success routes. If we filter out unrealistic reactions using log-likelihood under the reference backward model, we can improve the performance of retrosynthetic planning while maintaining that of the backward reaction model. We report mean and standard deviation across five independent runs.

FILTERING THRESHOLD ϵ	REACTIONS		REACTION PATHWAYS				
	TOP-1 \uparrow	TOP-10 \uparrow	SUCC. RATE \uparrow ($N = 50$)	SUCC. RATE \uparrow ($N = 500$)	LENGTH \downarrow	TIME \downarrow	COST \downarrow
0.9	44.45 \pm 0.01	73.29 \pm 0.01	48.42 \pm 0.88	92.21 \pm 0.70	8.77 \pm 0.11	130.66 \pm 1.69	13.23 \pm 0.40
0.8	44.52 \pm 0.01	73.30 \pm 0.01	47.68 \pm 1.23	92.00 \pm 0.21	8.74 \pm 0.08	129.70 \pm 2.15	13.01 \pm 0.18
0.7	44.50 \pm 0.02	73.28 \pm 0.01	45.79 \pm 0.58	91.47 \pm 0.84	8.82 \pm 0.21	131.91 \pm 1.84	13.15 \pm 0.38
0.6	44.48 \pm 0.01	73.28 \pm 0.01	44.11 \pm 0.52	90.95 \pm 1.02	8.92 \pm 0.28	130.14 \pm 1.02	13.55 \pm 0.62
0.5	44.46 \pm 0.01	73.26 \pm 0.00	43.79 \pm 0.61	91.05 \pm 0.67	8.95 \pm 0.20	130.39 \pm 2.69	13.24 \pm 0.43
0	41.04 \pm 0.02	71.94 \pm 0.01	54.32 \pm 0.39	83.58 \pm 0.39	9.91 \pm 0.09	149.16 \pm 1.90	18.23 \pm 0.18

et al., 2018), DFPN-E (Kishimoto et al., 2019), RETRO*, and RETRO*-0 (Chen et al., 2020b). We note that the considered baselines focus on designing efficient search algorithms for retrosynthetic planning, while our main contribution is the training of the backward reaction model towards maximizing the performance of retrosynthetic planning.

4.2. Main Result

We report the results for evaluating the performance of our framework and baselines in Table 1 and Figure 4. We observe that combining our framework with RETRO* and RETRO*-0 significantly outperforms the baselines, including the original RETRO* and RETRO*-0. For example, RETRO*-0+OURS achieves the success rate of 96.32% with a computation limit of $N = 500$, while RETRO* achieves 86.84% as the best baseline. We also observe that RETRO*-0+OURS and RETRO*+OURS outperforms RETRO*-0 and

RETRO* significantly in terms of other evaluation metrics, e.g., the length and the cost of discovered reaction pathways.⁴ Such a result demonstrates the effectiveness of our framework for retrosynthetic planning.

On the other hand, our backward reaction model does not suffer from a drop in TOP- k accuracy. Instead, they even outperform the backward reaction model trained on the reaction dataset $\mathcal{D}_{\text{reaction}}$ via supervised learning in terms of TOP-10 accuracy. We understand that such an improvement in the TOP-10 accuracy comes from "diverse" solution candidates generated by our backward reaction model, which is encouraged by being trained on a large variety of samples, e.g., augmented reactions. Intriguingly, we also observe that

⁴If a reaction pathway is failed to be found, the length and the cost are set to two times of maximum length and cost among the ground-truth pathways of $\mathcal{D}_{\text{target}}$, respectively, and the time is set to the limit of backward reaction model calls, i.e., 500.




Figure 6. Reaction pathways from the same target molecule, searched by (a) RETRO*-0 + OURS and (b) RETRO*-0. In the example, our reaction pathway has a shorter length, which implies that our solution has better quality than that from RETRO*-0, as shorter reaction pathways are easier to be conducted in laboratories.

RETRO*-0+OURS performs similarly with RETRO*+OURS without the help of an additional value function for guiding the search algorithm, i.e., RETRO*. We hypothesize that the performance of the search algorithm may saturate when the backward reaction model appropriately adapts to the search algorithm. This highlights the importance of training an appropriate backward reaction model for retrosynthetic planning.

4.3. Ablation Study

Next, we conduct ablation studies on our framework to investigate the effect of (1) removing the reaction augmentation procedure, (2) varying the number of iterations, and (3) varying the filtering threshold ϵ for realistic reactions. For the ablation studies on (1) and (3), our overall procedure is conducted a single time, i.e., we do not iterate our framework multiple times.

Effectiveness of reaction augmentation. To evaluate whether if the proposed reaction augmentation is effective, we evaluate our framework without the reaction augmentation scheme. As shown in Table 2, our framework achieves a higher success rate for finding the reaction pathway for the target molecules. This validates the effectiveness of our reaction augmentation scheme based on the forward reaction model.

Number of iterations. In Figure 5, we report the perfor-

mance of the backward reaction model and the success rate for finding the reaction pathways over multiple iterations. The result demonstrates that iterating our framework improves the success rate of finding a reaction pathway while maintaining the accuracy of the backward reaction model. This indeed validates that our framework is effective without compromise in the ability to model realistic reactions.

Filtering threshold ϵ . To recognize the effectiveness of filtering out unrealistic reactions within success routes, we compare the performance of our framework with varying thresholds for filtering out reactions.⁵ As shown in Table 3, the performance of our backward reaction model is (a) improved when using the threshold, i.e., using $\text{Thr}>0$ instead of $\text{Thr}=0$, and (b) robust to the change of hyper-parameter, i.e., varying the threshold from $\{0.5, 0.6, 0.7, 0.8, 0.9\}$. In particular, setting FILTERING THRESHOLD ϵ to 0, the top-1 exact match accuracy of our backward reaction model degrades from 44.52% to 41.04%. On the other hand, if we filter out unrealistic reactions by the reference backward reaction model, our backward reaction model retains its reliability and enjoys the improved performance in retrosynthetic planning without compromise.

⁵In this experiment, we do not include the augmented reaction data via the forward reaction model.

4.4. Case Study

In Figure 6, we compare reaction pathways found by RETRO*-0+OURS and RETRO*-0 for the same target molecule. In the example, we observe that both frameworks can find realistic reaction pathways for the given target molecule. However, the reaction pathway searched by RETRO*-0+OURS is preferable according to our predefined criteria, i.e., the number of reactions required for execution.

5. Conclusion

We propose a new framework based on self-improved model adaptation to improve retrosynthetic planning. Our main idea is to train a backward reaction model to imitate success routes found by the retrosynthetic planning algorithm combined with itself. We also propose an additional augmentation scheme that diversifies training through generating new reactions from a forward reaction model. Experiments show that our framework successfully adapts the backward reaction model to generate reactions that are both realistic and executable from building block molecules. Meanwhile, improving filtering modules using negative (unrealistic) reaction samples and developing better reaction augmentation schemes would be interesting directions to explore.

Acknowledgements

We thank Yeonghun Kang, Seonyul Kim, and Sangwoo Mo for providing helpful feedback and suggestions in preparing the early version of the manuscript. We would like to thank Binghong Chen for providing the dataset and source implementation of RETRO*. This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)). This work was supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2019-0-01396, Development of framework for analyzing, detecting, mitigating of bias in AI model and training data).

References

- Anthony, T., Tian, Z., and Barber, D. Thinking fast and slow with deep learning and tree search. *arXiv preprint arXiv:1705.08439*, 2017.
- Anthony, T., Nishihara, R., Moritz, P., Salimans, T., and Schulman, J. Policy gradient search: Online planning and expert iteration without search trees. *arXiv preprint arXiv:1904.03646*, 2019.
- Bradshaw, J., Kusner, M. J., Paige, B., Segler, M. H., and Hernández-Lobato, J. M. A generative model for electron paths. *arXiv preprint arXiv:1805.10970*, 2018.
- Chen, B., Dai, B., Lin, Q., Ye, G., Liu, H., and Song, L. Learning to plan in high dimensions via neural exploration-exploitation trees. In *International Conference on Learning Representations*, 2020a.
- Chen, B., Li, C., Dai, H., and Song, L. Retro*: learning retrosynthetic planning with neural guided a* search. In *International Conference on Machine Learning*, pp. 1608–1616. PMLR, 2020b.
- Coley, C. W., Green, W. H., and Jensen, K. F. Rdkchiral: An rdkit wrapper for handling stereochemistry in retrosynthetic template extraction and application. *Journal of chemical information and modeling*, 59(6):2529–2537, 2019.
- Corey, E. J. The logic of chemical synthesis: multistep synthesis of complex carbogenic molecules (nobel lecture). *Angewandte Chemie International Edition in English*, 30(5):455–465, 1991.
- Dai, H., Li, C., Coley, C., Dai, B., and Song, L. Retrosynthesis prediction with conditional graph logic network. In *Advances in Neural Information Processing Systems*, pp. 8872–8882, 2019.
- Do, K., Tran, T., and Venkatesh, S. Graph transformation policy network for chemical reaction prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 750–760, 2019.
- Hughes, J. P., Rees, S., Kalindjian, S. B., and Philpott, K. L. Principles of early drug discovery. *British journal of pharmacology*, 162(6):1239–1249, 2011.
- Jin, W., Coley, C. W., Barzilay, R., and Jaakkola, T. Predicting organic reaction outcomes with weisfeiler-lehman network. *arXiv preprint arXiv:1709.04555*, 2017.
- Karpov, P., Godin, G., and Tetko, I. V. A transformer model for retrosynthesis. In *International Conference on Artificial Neural Networks*, pp. 817–830. Springer, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kishimoto, A., Buesser, B., Chen, B., and Botea, A. Depth-first proof-number search with heuristic edge cost and application to chemical synthesis planning. In *Advances in Neural Information Processing Systems*, pp. 7226–7236, 2019.
- Liu, B., Ramsundar, B., Kawthekar, P., Shi, J., Gomes, J., Luu Nguyen, Q., Ho, S., Sloane, J., Wender, P., and

- Pande, V. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS central science*, 3(10):1103–1113, 2017.
- Liu, C.-H., Korablyov, M., Jastrzębski, S., Włodarczyk-Pruszyński, P., Bengio, Y., and Segler, M. H. Retrognn: Approximating retrosynthesis by graph neural networks for de novo drug design. *arXiv preprint arXiv:2011.13042*, 2020.
- Lowe, D. M. *Extraction of chemical structures and reactions from the literature*. PhD thesis, University of Cambridge, 2012.
- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Schreck, J. S., Coley, C. W., and Bishop, K. J. Learning retrosynthetic planning through simulated experience. *ACS central science*, 5(6):970–981, 2019.
- Schwaller, P., Gaudin, T., Lanyi, D., Bekas, C., and Laino, T. “found in translation”: predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models. *Chemical science*, 9(28):6091–6098, 2018.
- Schwaller, P., Laino, T., Gaudin, T., Bolgar, P., Hunter, C. A., Bekas, C., and Lee, A. A. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS central science*, 5(9):1572–1583, 2019.
- Segler, M. H. and Waller, M. P. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry—A European Journal*, 23(25):5966–5971, 2017.
- Segler, M. H., Preuss, M., and Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018.
- Shi, C., Xu, M., Guo, H., Zhang, M., and Tang, J. A graph to graphs framework for retrosynthesis prediction. *arXiv preprint arXiv:2003.12725*, 2020.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Somnath, V. R., Bunne, C., Coley, C. W., Krause, A., and Barzilay, R. Learning graph models for template-free retrosynthesis. *arXiv preprint arXiv:2006.07038*, 2020.
- Song, J., Lanka, R., Zhao, A., Bhatnagar, A., Yue, Y., and Ono, M. Learning to search via retrospective imitation. *arXiv preprint arXiv:1804.00846*, 2018.
- Yan, C., Barlow, S., Wang, Z., Yan, H., Jen, A. K.-Y., Marder, S. R., and Zhan, X. Non-fullerene acceptors for organic solar cells. *Nature Reviews Materials*, 3(3):1–19, 2018.
- Yan, C., Ding, Q., Zhao, P., Zheng, S., Yang, J., Yu, Y., and Huang, J. Retroxpert: Decompose retrosynthesis prediction like a chemist. In *Advances in Neural Information Processing Systems*, 2020.
- Zheng, S., Rao, J., Zhang, Z., Xu, J., and Yang, Y. Predicting retrosynthetic reactions using self-corrected transformer neural networks. *Journal of Chemical Information and Modeling*, 60(1):47–55, 2019.