

---

# DirectMultiStep: Direct Route Generation for Multi-Step Retrosynthesis

---

**Yu Shee\***  
Yale University  
yu.shee@yale.edu

**Haote Li\***  
Yale University  
haote.li@yale.edu

**Anton Morgunov\***  
Yale University  
anton.morgunov@yale.edu

**Victor Batista**  
Yale University  
victor.batista@yale.edu

## Abstract

Traditional computer-aided synthesis planning (CASP) methods rely on iterative single-step predictions, leading to exponential search space growth that limits efficiency and scalability. We introduce a transformer-based model that directly generates multi-step synthetic routes as a single string by conditionally predicting each molecule based on all preceding ones. The model accommodates specific conditions such as the desired number of steps and starting materials, outperforming state-of-the-art methods on the PaRoutes dataset with a 2.2x improvement in Top-1 accuracy on the  $n_1$  test set and a 3.3x improvement on the  $n_5$  test set. It also successfully predicts routes for FDA-approved drugs not included in the training data, showcasing its generalization capabilities. While the current suboptimal diversity of the training set may impact performance on less common reaction types, our approach presents a promising direction towards fully automated retrosynthetic planning.

## 1 Introduction

Finding the most efficient route to a desired chemical compound is a daily challenge for synthetic organic chemists in both fundamental research and drug discovery. Route efficiency is determined by various factors, some of which can be objectively assessed, such as overall yield (not all chemical reactions have 100% conversion rate) and enantiomeric excess (in case of chiral compounds), where higher values are always preferred. Other factors, such as atom efficiency (minimization of byproducts) and availability (cost) of starting materials, are more case-dependent. A used chemical reactant is considered waste unless it can be repurposed as a reactant in a different process. Similarly, the choice of starting materials depends on factors such as budget, logistics, and the availability of specific equipment. It's worth noting that many commercially available compounds can be synthesized from other commercially available compounds, adding another layer of complexity to the decision-making process.

A significant contribution to the development of algorithmic frameworks for identifying synthetic routes leading to any desired compound (subject of the 1990 Nobel Prize in Chemistry) was made by Elias James Corey. Corey's framework, now known as retrosynthetic analysis, begins with identifying atoms that would serve as reaction centers. Disconnecting bonds between these centers results in the formation of hypothetical fragments (called synthons) from which a precursor molecule can be created. This mapping from fragments to actual molecules is one-to-many because there is usually

---

\*Equal contribution. Listing order is random.

more than one functional group that could participate in any given type of reaction. Importantly, meticulous application of Corey’s framework (i.e., systematically breaking small subsets of bonds) will eventually lead to commercially available starting materials. The algorithmic nature of this process allowed Corey to envision automating these rules to create Computer-Aided Synthesis Planning (CASP) as early as 1969 [1].

Recent advancements in data science and machine learning (ML) methods have led to a surge of interest in developing CASP methods [2, 3]. The vast majority of existing methods [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35] are designed to automate single-step retrosynthetic (SSR) analysis, i.e. predicting a list of compounds from which a target product could be made in one step. A full multi-step route can be determined by iteratively applying SSR methods to the identified precursors until a termination condition is satisfied (e.g., identifying reactants in the database of commercially available compounds). Notably, because each SSR method call creates a list of candidates, iterative application of these methods generates an exponentially growing search space. Graph traversal algorithms such as Monte Carlo Tree Search (MCTS), Depth-First Proof Number (DFPN) search, and A-star (Retro\*) have been applied to efficiently traverse this exponential search space [36, 37, 38].

However, evaluation of performance of these SSR methods on full route prediction has been limited by the scarcity of open-source datasets containing valid multi-step routes. The creation of PaRoutes [39], an open-source dataset containing 450k multi-step routes extracted from the United States Patent and Trademark Office (USPTO), marked a major development. Notably, even state-of-the-art SSR models [6, 40] combined with advanced search algorithms (MCTS, DFPN, Retro\*) correctly identify multi-step routes (Top-1 accuracy) for only 17% and 10% of the target compounds in  $n_1$  and  $n_5$ , respectively. This performance can be rationalized by recognizing that retrosynthesis is inherently a multi-step problem: the optimal choice of reaction to make a compound depends on subsequent steps in the synthesis. For example, a common pattern in multi-step routes includes: 1) protection of certain functional groups, 2) the desired transformation of unprotected groups, and 3) deprotection of the protected groups. An SSR method applied to the reactant of reaction 2 may output numerous candidate precursors; however, knowing that protective groups are removed in subsequent steps dramatically changes the probability distribution over those precursors.

In this work, we propose a novel approach for direct prediction of multi-step routes, bypassing the need for single-step models and sophisticated exponential graph traversal algorithms. Our DirectMultiStep model demonstrates state-of-the-art performance on  $n_1$  and  $n_5$  evaluation sets and shows generalizability by successfully predicting routes for FDA-approved drugs absent from the training set. We also discuss the model’s limitations, which may include the need to split longer routes into sub-routes, provide the structure of the starting material, or supply a precursor to the target compound when dealing with transformations not well-represented in the training dataset. Despite these limitations, we believe our results are promising enough to warrant further attention to developing retrosynthetic methods that directly predict multi-step routes.

## 2 Methods

### 2.1 Definitions and multi-step route representation

This subsection defines key terms related to our multi-step retrosynthetic methodology. The "target compound" (green in Fig. 1) represents the final product of the multi-step synthesis tree. Starting materials (SM, orange in Fig. 1) are compounds for which no further precursors need to be identified, i.e., they are the leaves of the synthesis tree. The "number of steps" denotes the largest number of reactions from SM to the target compound (tree height). In our approach, routes are represented as recursive dictionaries (Fig. 1a) containing the SMILES representation of the molecule and a list of other dictionaries, containing either starting materials or trees leading to precursors of the current node. Removing space and newline characters creates a string representation predicted by our transformer model. We train the DirectMultiStep (DMS) transformer, which takes the target compound, (optionally) starting material, and number of steps as input, and predicts the multi-step route represented as a string.

## 2.2 Data curation and preparation

A second version of the PaRoutes [39] dataset containing 450k routes is stripped of all metadata and stored as recursive dictionaries representing multi-step routes. The two evaluation sets,  $n_1$  and  $n_5$ , are processed similarly. A training partition is created by removing all permutations (swapping left and right subtrees) of routes in  $n_1$  and  $n_5$  from the full dataset, resulting in 407k routes. The training dataset is augmented by adding 2 permutations for each route. To train the DMS model which takes starting material as an input, we find all starting materials (leaves) for each tree and store a combination of the target compound with each starting material as a separate entry. As a result, DMS with SM is trained on 3 349 118 inputs, and the DMS without SM on 1 078 951 inputs. All SMILES strings are tokenized by treating each character as a token, and the string representation of the multi-step route is tokenized similarly, treating delimiters of the tree (smiles', children', [, ], {, and }) as separate tokens. The final vocabulary size is 52 (including start, end, and padding tokens), the largest multi-step route has 1074 tokens, while the largest target compound and starting material have 145 and 135 tokens, respectively.




Figure 1: The workflow of DirectMultiStep. (a) The SMILES representation of the target compound (green), (optional) starting material (orange), and the number of steps are tokenized, concatenated, and fed into our transformer model. The model predicts a string representation of the multi-step synthesis tree. Spaces are added for clarity, and indentations indicate the levels in the synthesis route (tree). (b) Molecular structures corresponding to the target compound (green), (optional) starting material (orange), and the predicted synthesis tree with structures of all molecules.

## 2.3 Model architecture and training

We present two variants of the DirectMultiStep (DMS) model: DMS-10M and DMS-60M, with approximately 10 million and 60 million parameters, respectively. Both models employ the transformer encoder-decoder architecture [41]. DMS-10M has 6 layers and 8 attention heads in both encoder and decoder, with inputs embedded in 256 dimensions and amplified 3x in the feedforward blocks. DMS-60M has 8 layers and 8 attention heads, with inputs embedded in 512 dimensions and amplified 4x in the feedforward blocks. Both models use GeLU activations [42]. The encoder of DMS with SM receives a concatenation of the tokenized target compound, starting material, and number of steps, while the encoder of DMS without SM receives only the tokenized target compound and number of steps. In all models, the decoder receives output from the encoder and the start token of the multi-step route.

During training, a 10% dropout rate is used in embedding, attention, and feedforward layers. The




Figure 2: Distribution of the relative frequencies of route lengths (in terms of number of steps) in the training dataset before augmentation with permutations (407 148 routes, black),  $n_1$  test set (10 000 routes, blue), and  $n_5$  test set (10 000 routes, purple). Distribution is split into routes shorter (left subplot) and longer than 6 steps (right subplot).

ADAM optimizer [43] is used with a learning rate scheduler that warms up from 0 to  $3 \times 10^{-4}$  for DMS-10M and  $1 \times 10^{-4}$  for DMS-60M during the first 10% of training steps and undergoes cosine decay to  $3 \times 10^{-5}$  over 20 epochs. Gradients are clipped to 1.0. Finally, during training, each token in the input to the encoders has a 5% chance of being masked. Training requires approximately 200 clock hours using 4 NVIDIA A100-40 GB GPUs on an external cluster. An additional 450 hours were used for experimentation with model architectures and optimization of hyperparameters.

## 2.4 Beam search for route generation

We utilize beam search with a width of 50 to generate several candidates for the same target compound (and starting material, if provided). The cumulative probability of each beam serves as the ranking metric for prediction accuracy (Top- $K$ ). After beam search, predicted routes are checked for the validity of all SMILES strings (which are canonicalized), the correctness of the target compound, the presence of starting materials in the stock compound (if applicable), and the absence of repetitions.

## 2.5 Evaluation metrics and post-processing

Evaluation is performed on all 10 000 routes in both  $n_1$  and  $n_5$ , with the deepest leaf selected as the starting material to be included in the input to the encoder (for models that use starting materials). Top- $K$  accuracy evaluates the percentage of top  $K$  predictions precisely matching correct multi-step routes from the test set. Predictions matching any permutation of the correct route are counted as matches.

# 3 Results and Discussion

## 3.1 Test accuracy

Top- $K$  accuracy is shown in Table 1 for set- $n_1$  and Table 2 for set- $n_5$ . Our DirectMultiStep-SM-10M model achieves state-of-the-art performance, surpassing all previous methods in all Top- $K$  accuracies, with the most significant increase in Top-1 accuracy: 2.2x on  $n_1$  and 3.3x on  $n_5$ . Remarkably, our model’s Top-3 accuracy exceeds the Top-10 accuracy of top tree search methods using single-step retrosynthesis (SSR), highlighting the advantages of the multi-step first approach.

The DirectMultiStep-SM-60M model, despite having six times more parameters than the 10M variant, exhibits lower performance on higher Top- $K$  accuracies, suggesting that the 10M model size is optimal for the current dataset when starting material structure is provided. However, achieving comparable performance without starting material information requires a larger model, as demonstrated by the DirectMultiStep-noSM-60M model’s performance. These findings reveal a trade-off between model generality and size at the current dataset scale: less input information necessitates a larger model to maintain high performance. In all subsequent analysis, we use the routes predicted by the DirectMultiStep-SM-10M model.

Table 1: Top- $K$  accuracy on route test set- $n_1$  (10 000 routes).

Method	Top-1	Top-2	Top-3	Top-4	Top-5	Top-10
MCTS <sup>a</sup>	0.17				0.46	0.49
Retro* <sup>a</sup>	0.15				0.41	0.45
DFPN <sup>a</sup>	0.11				0.17	0.17
DirectMultiStep-SM-10M	<b>0.38</b>	0.46	0.50	0.52	<b>0.53</b>	<b>0.56</b>
DirectMultiStep-SM-60M	<b>0.38</b>	0.46	0.48	0.50	0.51	0.52
DirectMultiStep-noSM-60M	0.36	0.42	0.44	0.45	0.45	0.46

<sup>a</sup>These data are collected from the 2.0 version of PaRoutes in their GitHub repository [39].

Table 2: Top- $K$  accuracy on route test set- $n_5$  (10 000 routes).

Method	Top-1	Top-2	Top-3	Top-4	Top-5	Top-10
MCTS <sup>a</sup>	0.10				0.28	0.33
Retro* <sup>a</sup>	0.10				0.30	0.36
DFPN <sup>a</sup>	0.05				0.07	0.07
DirectMultiStep-SM-10M	0.33	0.40	0.43	0.44	<b>0.46</b>	<b>0.48</b>
DirectMultiStep-SM-60M	<b>0.34</b>	0.41	0.43	0.44	0.45	0.46
DirectMultiStep-noSM-60M	0.32	0.37	0.39	0.40	0.41	0.41

<sup>a</sup>These data are collected from the 2.0 version of PaRoutes in their GitHub repository [39].

Importantly, our model requires only one call to obtain multiple routes for each target compound, instead of hundreds of single-step model calls and even more reaction template applications (as shown in PaRoutes [39]). Furthermore, our model’s inference time (beam size 50) on a single GPU is comparable (3-16 seconds depending on the number of steps) to the time required to find the first successful route using previous methods (7-50 seconds [39]).

Fig. 3 shows the distribution of Top-1 and Top-10 accuracy over different route lengths on both  $n_1$  and  $n_5$ . Given that 90% of routes in the training partition have 4 or fewer steps Fig. 2), one would expect accuracy to decrease dramatically with increasing route length. However, the performance on routes with 5-8 steps is comparable to that of shorter routes, and the performance on 9-step routes is even comparable to that of 2-step routes. The exceptional accuracy on 9-step routes could be rationalized by the fact that their relative abundance in  $n_1$  and  $n_5$  does not exceed the relative abundance in the training partition (unlike for routes with length 3-7, as seen in Fig. 3).




Figure 3: Distribution of Top-1 and Top-10 accuracy of predictions with our model on test sets  $n_1$  and  $n_5$ . There is only one route with length 10 in  $n_1$ , and our model does not predict it correctly. That route is reproduced by splitting it in half, as shown in Fig. 4

The requirement to provide the starting material structure as input to the model is a limitation, although not as significant as it may seem. Corey’s retrosynthetic framework allows finding at least one route to starting materials by performing one functional group transformation at a time, usually requiring protection of other functional groups. However, such a route will be too long to be efficient in terms of yield (overall yield is the product of yields of individual steps). The art of

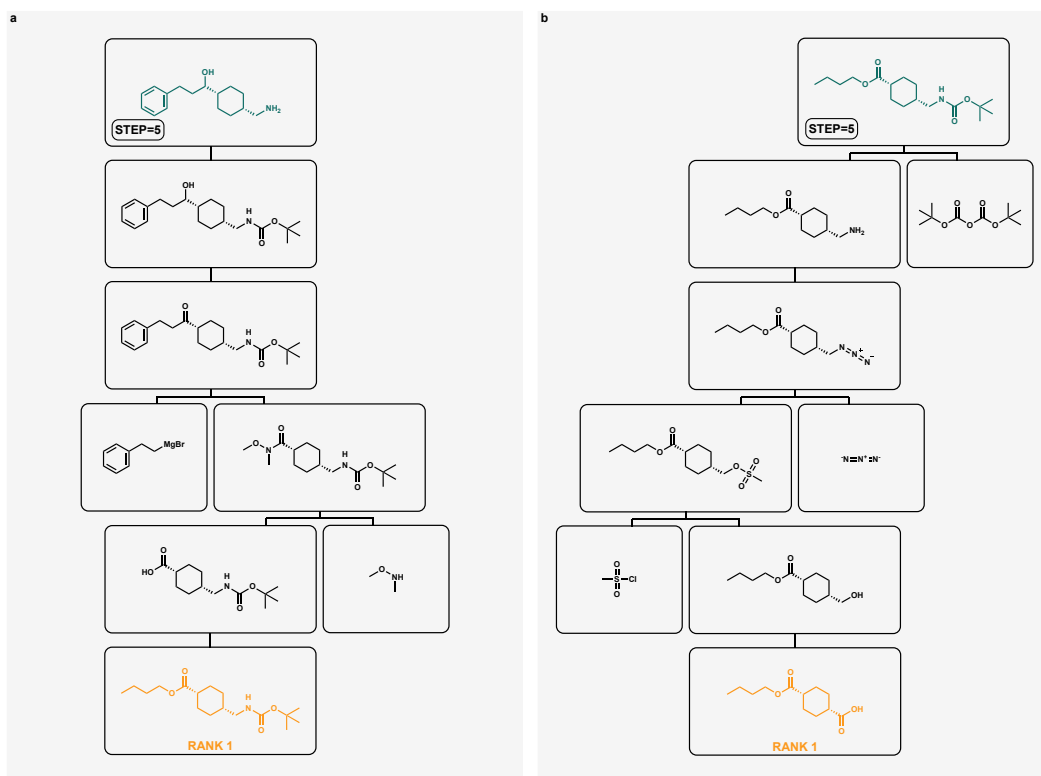


Figure 4: Separation of a 10-step route from set- $n_1$  into two 5-step routes. (a) Correct prediction for the first half of the 10-step route with starting material information. (b) Correct prediction for the second half of the 10-step route with starting material information.

organic synthesis lies in finding ways to perform several transformations simultaneously using a minimal number of protecting groups. We believe our model provides value as an assistant to this task: finding better ways to get to the target compound from a starting material. The requirement to specify the number of steps as an input becomes a useful option in this context. Moreover, if the target compound is chiral (i.e., it has a mirror image which cannot be superimposed), even during manual retrosynthetic planning, chemists check if chiral centers can be obtained from chiral starting materials (such as naturally occurring amino acids) because developing reactions that create chiral centers in the desired conformation remains an active research area (and subject of 2001 and 2021 Nobel Prizes in Chemistry). An example of such a route is shown in Fig. 6.

### 3.2 Separation of Long Routes into Shorter Routes

As seen in Fig. 3, our model does not find the single 10-step route in the  $n_1$  set. However, this route can be correctly reproduced with rank 1 if split into two 5-step routes and the intermediate compound is provided as the starting material for the first 5 reactions (Fig. 4a) and a target compound for the last 5 reactions (Fig. 4b). While ideally the necessity to provide an intermediate compound should be avoided, this resembles the bidirectional search technique commonly employed in manual retrosynthetic analysis [44].

### 3.3 Retrosynthetic planning of pharmaceutical compounds

We demonstrate the generalizability of our model by testing on three FDA-approved drugs: Vonoprazan, Mitapivat, and Daridorexant, which were used for evaluation by Xiong, et al. [45]. These drugs and their intermediates are absent from the training set. Vonoprazan, a potassium competitive acid blocker for *Helicobacter pylori* infections, was initially proposed with a 2-step route (Fig. 5a) [46]. However, one of its starting materials is unstable, and the final reaction leads to significant byproducts. Subsequently, an improved 4-step route (Fig. 5b) was introduced [47]. Mitapivat, a pyruvate kinase

activator for treating hemolytic anemia, has been associated with both a 5-step route (Fig. 5c) and a 3-step route (Fig. 5d) [48, 49]. Daridorexant, an orexin receptor antagonist for adult insomnia, is linked to a 4-step route (Fig. 6) [50].

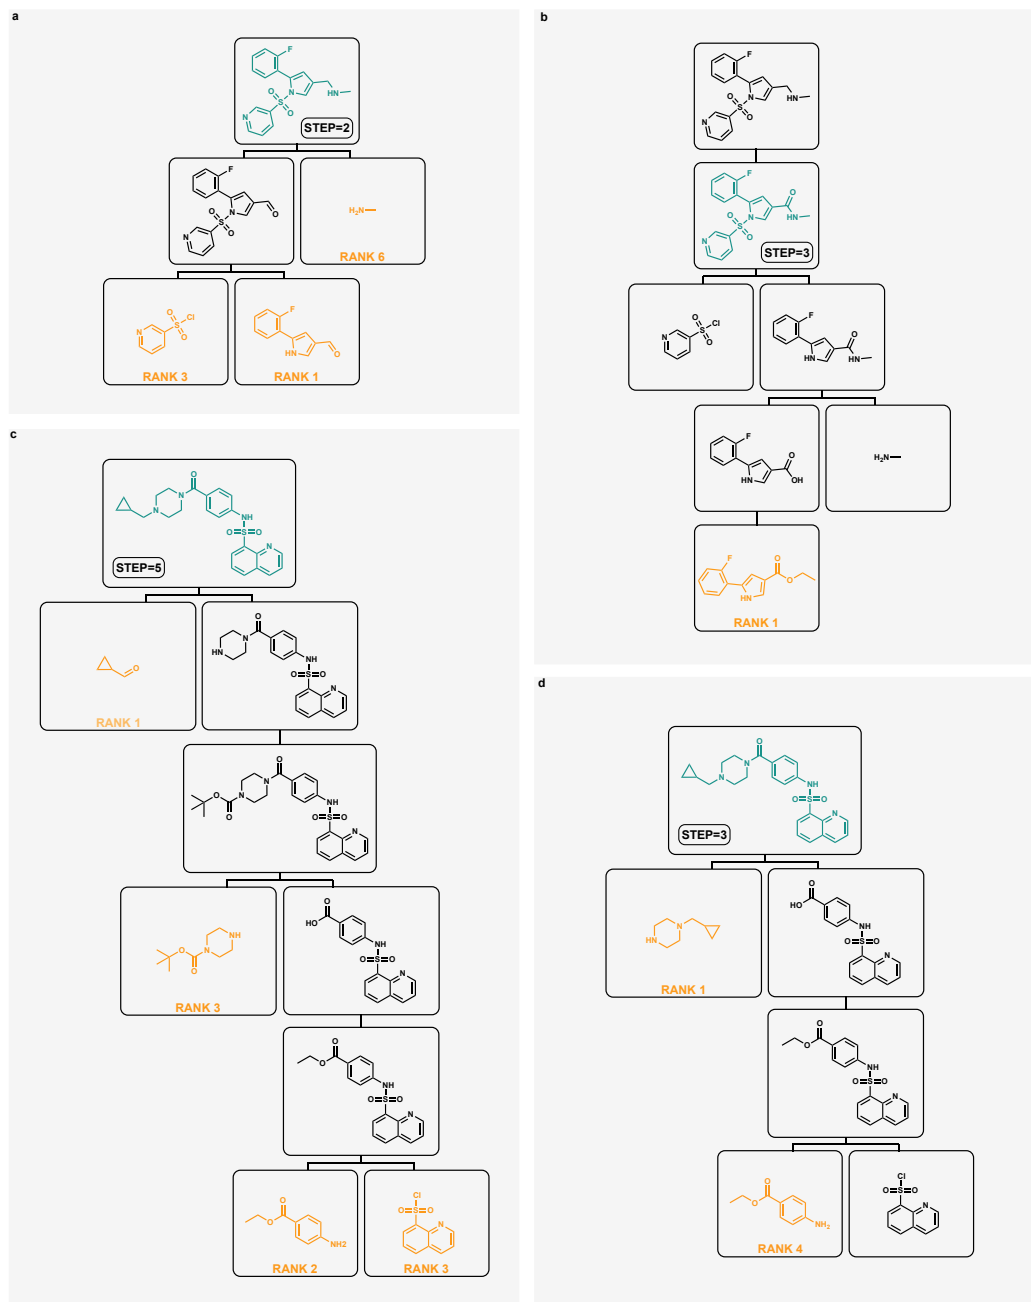


Figure 5: Literature routes for Vonoprazan and Mitapivat correctly reproduced by our model. Target compounds are in green, starting materials that are given as inputs are colored in orange. Ranks denote the rank of this route when the specified starting material is provided. (a) First literature route for Vonoprazan. The model predicts correctly no matter which starting material is given. (b) Second literature route for Vonoprazan. The model predicts the route correctly only when an immediate precursor to Vonoprazan is given as the target compound. (c) First literature route for Mitapivat. (d) Second literature route for Mitapivat.

Our model correctly reproduces Vonoprazan's first literature route (Fig. 5a) regardless of the starting material provided, with high ranks. However, the model struggles to find the first transformation (counting from the target compound) from the second route. In the first route, the first step is a reductive amination, while in the second route, it's a reduction of the carbonyl bond in the amide group. Reductive amination is much more common for creating C-N bonds, so the training set is heavily biased toward it. As a result, all routes predicted by our model start as in Fig. 5a. However, if we provide the precursor to Vonoprazan in the second route as the target compound, the model correctly reproduces the full route (Fig. 5b). This highlights an important limitation of our model: it may not perform as well on routes involving transformations that are inadequately represented in the training dataset. We envision that as larger multi-step datasets become available, this problem will naturally disappear.




Figure 6: Daridorexant literature multi-step routes and predictions from the model with starting material information. The compounds that are missing from the prediction are in red.

Our model correctly reproduces both routes for Mitapivat (Fig. 5c and Fig. 5d). In contrast, the route for Daridorexant (Fig. 6) is only partially predicted: our model does not split the right child of the root node (Daridorexant itself) into the two compounds highlighted in red. All other intermediates and starting materials are predicted correctly. The reason for this behavior is not well understood, but it can be rationalized by looking at the distribution of the number of leaves at root nodes: 73% of root nodes (target compounds) have at least 1 leaf (see SI Fig. 7). Such dependence on the distributions of the training set comprises an important limitation of the multi-step first approach. This issue would not occur with methods based on single-step prediction because the tree search framework would not halt at compounds outside the stock compound set before reaching the maximum iterations. Ensuring that multi-step first models reproduce that aspect of SSR methods constitutes a direction for future research.



## 4 Conclusion

In this work, we introduced DirectMultiStep, a transformer-based model that directly generates multi-step synthetic routes for a given target compound. By predicting the entire route as a single string, our approach bypasses the need for iterative single-step predictions and exponential search space traversal. The model accommodates specific conditions such as the desired number of steps and starting materials, enabling efficient route planning.

DirectMultiStep outperforms state-of-the-art methods on the PaRoutes dataset, achieving a 2.2x improvement in Top-1 accuracy on the  $n_1$  test set and a 3.3x improvement on the  $n_5$  test set. The model also successfully predicts routes for FDA-approved drugs not included in the training data, showcasing its generalization capabilities. However, the model’s performance may be impacted by the suboptimal diversity of the current training set, particularly for less common reaction types.

Despite these limitations, the DirectMultiStep approach presents a promising direction towards fully automated retrosynthetic planning. As larger and more diverse multi-step datasets become available, the model’s performance is expected to improve further. Future research should focus on addressing the limitations identified, such as the need to split longer routes into sub-routes and the dependence on the distributions of the training set. Overall, this work demonstrates the potential of transformer-based models for efficient and scalable computer-aided synthesis planning.

## 5 Broad Impact

Chemical synthesis of target compounds of interest forms the routine for millions of organic chemists around the world, including graduate students in PhD programs and synthetic chemists in pharmaceutical companies. The ability to find shorter (and thus more efficient) routes to any desired compound is an enticing challenge with the potential of saving on labor costs and chemical waste. As a result, multi-step prediction methods have the potential to improve the lives of millions of people, and potentially contribute to faster advances in the development of new drugs and materials. We hope that our work constitutes an important advancement in that direction.

Quite unfortunately, however, there is no fundamental chemical difference between how safe and useful compounds (such as drugs) and malign compounds (such as chemical weapons) are synthesized. Both require application of the same set of rules of organic synthesis, and so the ability to make drugs faster and more efficiently inherently brings along the ability to make poisons faster and more efficiently. One cannot impose restrictions on the tokens (chemical elements) learned by the model, as even some combinations of nitrogen and oxygen (which are widespread in proteins and DNA) may be cytotoxic. Therefore, unless the humanity at large decides that the benefits of automated retrosynthesis do not outweigh its potential harms, the most we can do is communicate the dangers clearly and perhaps impose restrictions that would still require a trained (which often accompanies training in ethics) chemist to use our model. Besides, all meta-data of reaction conditions and catalysts are removed from the predictions of our model, which would present a deliberate barrier against malicious uses.

## 6 Code and Data Availability

The reaction route dataset is from the 2.0 version of PaRoutes in their GitHub repository [39], available under the Apache License 2.0. Code to process dataset, implementation of model architecture, code for training, generation, and evaluation are available under MIT License at <https://github.com/batistagroup/DirectMultiStep>.

## 7 Acknowledgments

The authors acknowledge a generous allocation of high-performance computing time from NERSC. The development of the methodology was supported by the NSF CCI grant (VSB, Award Number 2124511).

## References

- [1] E. J. Corey and W. Todd Wipke. Computer-Assisted Design of Complex Organic Syntheses: Pathways for molecular synthesis can be devised with a computer and equipment for graphical communication. *Science*, 166(3902):178–192, 1969.
- [2] A. Filipa De Almeida, Rui Moreira, and Tiago Rodrigues. Synthetic organic chemistry driven by artificial intelligence. *Nat Rev Chem*, 3(10):589–604, 2019.
- [3] Thomas J. Struble, Juan C. Alvarez, Scott P. Brown, Milan Chytil, Justin Cisar, Renee L. DesJarlais, Ola Engkvist, Scott A. Frank, Daniel R. Greve, Daniel J. Griffin, Xinjun Hou, Jeffrey W. Johannes, Constantine Kreatsoulas, Brian Lahue, Miriam Mathea, Georg Mogk, Christos A. Nicolaou, Andrew D. Palmer, Daniel J. Price, Richard I. Robinson, Sebastian Salentin, Li Xing, Tommi Jaakkola, William. H. Green, Regina Barzilay, Connor W. Coley, and Klavs F. Jensen. Current and Future Roles of Artificial Intelligence in Medicinal Chemistry Synthesis. *J. Med. Chem.*, 63(16):8667–8682, 2020.
- [4] Zhongliang Guo, Stephen Wu, Mitsuru Ohno, and Ryo Yoshida. Bayesian Algorithm for Retrosynthesis. *J. Chem. Inf. Model.*, 60(10):4474–4486, October 2020.
- [5] Hankook Lee, Sungsoo Ahn, Seung-Woo Seo, You Young Song, Eunho Yang, Sung-Ju Hwang, and Jinwoo Shin. Retcl: A selection-based approach for retrosynthesis via contrastive learning, 2021.
- [6] Marwin H. S. Segler and Mark P. Waller. Neural-Symbolic Machine Learning for Retrosynthesis and Reaction Prediction. *Chem. Eur. J.*, 23(25):5966–5971, May 2017.
- [7] Connor W. Coley, Luke Rogers, William H. Green, and Klavs F. Jensen. Computer-Assisted Retrosynthesis Based on Molecular Similarity. *ACS Cent. Sci.*, 3(12):1237–1245, December 2017.
- [8] Shoichi Ishida, Kei Terayama, Ryosuke Kojima, Kiyosei Takasu, and Yasushi Okuno. Prediction and Interpretable Visualization of Retrosynthetic Reactions Using Graph Convolutional Networks. *J. Chem. Inf. Model.*, 59(12):5026–5033, December 2019.
- [9] Michael E. Fortunato, Connor W. Coley, Brian C. Barnes, and Klavs F. Jensen. Data Augmentation and Pretraining for Template-Based Retrosynthetic Prediction in Computer-Aided Synthesis Planning. *J. Chem. Inf. Model.*, 60(7):3398–3407, July 2020.
- [10] Hanjun Dai, Chengtao Li, Connor W. Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network, 2020.
- [11] Shuan Chen and Yousung Jung. Deep Retrosynthetic Reaction Prediction using Local Reactivity and Global Attention. *JACS Au*, 1(10):1612–1620, October 2021.
- [12] Philipp Seidl, Philipp Renz, Natalia Dyubankova, Paulo Neves, Jonas Verhoeven, Jörg K. Wegner, Marwin Segler, Sepp Hochreiter, and Günter Klambauer. Improving Few- and Zero-Shot Reaction Template Prediction Using Modern Hopfield Networks. *J. Chem. Inf. Model.*, 62(9):2111–2120, May 2022.
- [13] Chaochao Yan, Qianggang Ding, Peilin Zhao, Shuangjia Zheng, Jinyu Yang, Yang Yu, and Junzhou Huang. Retroxpert: Decompose retrosynthesis prediction like a chemist, 2020.
- [14] Chence Shi, Minkai Xu, Hongyu Guo, Ming Zhang, and Jian Tang. A graph to graphs framework for retrosynthesis prediction, 2021.
- [15] Vignesh Ram Somnath, Charlotte Bunne, Connor W. Coley, Andreas Krause, and Regina Barzilay. Learning graph models for retrosynthesis prediction, 2021.
- [16] Xiaorui Wang, Yuquan Li, Jiezhong Qiu, Guangyong Chen, Huanxiang Liu, Benben Liao, Chang-Yu Hsieh, and Xiaojun Yao. RetroPrime: A Diverse, plausible and Transformer-based method for Single-Step retrosynthesis predictions. *Chemical Engineering Journal*, 420:129845, September 2021.
- [17] Yu Wang, Chao Pang, Yuzhe Wang, Junru Jin, Jingjie Zhang, Xiangxiang Zeng, Ran Su, Quan Zou, and Leyi Wei. Retrosynthesis prediction with an interpretable deep-learning framework based on molecular assembly tasks. *Nat Commun*, 14(1):6155, October 2023.
- [18] Weihe Zhong, Ziduo Yang, and Calvin Yu-Chian Chen. Retrosynthesis prediction using an end-to-end graph generative architecture for molecular graph editing. *Nat Commun*, 14(1):3009, May 2023.
- [19] Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic Reaction Prediction Using Neural Sequence-to-Sequence Models. *ACS Cent. Sci.*, 3(10):1103–1113, October 2017.
- [20] Pavel Karpov, Guillaume Godin, and Igor Tetko. A Transformer Model for Retrosynthesis. preprint, Chemistry, May 2019.
- [21] Benson Chen, Tianxiao Shen, Tommi S. Jaakkola, and Regina Barzilay. Learning to make generalizable and diverse predictions for retrosynthesis, 2019.
- [22] Alpha A. Lee, Qingyi Yang, Vishnu Sresht, Peter Bolgar, Xinjun Hou, Jacquelyn L. Klug-McLeod, and Christopher R. Butler. Molecular Transformer unifies reaction prediction and retrosynthesis across pharma chemical space. *Chem. Commun.*, 55(81):12152–12155, 2019.

- [23] Kangjie Lin, Youjun Xu, Jianfeng Pei, and Luhua Lai. Automatic retrosynthetic route planning using template-free models. *Chem. Sci.*, 11(12):3355–3364, 2020.
- [24] Shuangjia Zheng, Jiahua Rao, Zhongyue Zhang, Jun Xu, and Yuedong Yang. Predicting Retrosynthetic Reactions Using Self-Corrected Transformer Neural Networks. *J. Chem. Inf. Model.*, 60(1):47–55, January 2020.
- [25] Igor V. Tetko, Pavel Karpov, Ruud Van Deursen, and Guillaume Godin. State-of-the-art augmented NLP transformer models for direct and single-step retrosynthesis. *Nat Commun*, 11(1):5575, November 2020.
- [26] Seung-Woo Seo, You Young Song, June Yong Yang, Seohui Bae, Hankook Lee, Jinwoo Shin, Sung Ju Hwang, and Eunho Yang. GTA: Graph Truncated Attention for Retrosynthesis. *AAAI*, 35(1):531–539, May 2021.
- [27] Kelong Mao, Xi Xiao, Tingyang Xu, Yu Rong, Junzhou Huang, and Peilin Zhao. Molecular graph enhanced transformer for retrosynthesis prediction. *Neurocomputing*, 457:193–202, October 2021.
- [28] Mikołaj Sacha, Mikołaj Błaż, Piotr Byrski, Paweł Dąbrowski-Tumański, Mikołaj Chromiński, Rafał Loska, Paweł Włodarczyk-Pruszyński, and Stanisław Jastrzębski. Molecule Edit Graph Attention Network: Modeling Chemical Reactions as Sequences of Graph Edits. *J. Chem. Inf. Model.*, 61(7):3273–3284, July 2021.
- [29] Vipul Mann and Venkat Venkatasubramanian. Retrosynthesis prediction using grammar-based neural machine translation: An information-theoretic approach. *Computers & Chemical Engineering*, 155:107533, December 2021.
- [30] Umit V. Ucak, Taek Kang, Junsu Ko, and Juyong Lee. Substructure-based neural machine translation for retrosynthetic prediction. *J. Cheminform*, 13(1):4, December 2021.
- [31] Eunji Kim, Dongseon Lee, Youngchun Kwon, Min Sik Park, and Youn-Suk Choi. Valid, Plausible, and Diverse Retrosynthesis Using Tied Two-Way Transformers with Latent Variables. *J. Chem. Inf. Model.*, 61(1):123–133, January 2021.
- [32] Ross Irwin, Spyridon Dimitriadis, Jiazhen He, and Esben Jannik Bjerrum. Chemformer: a pre-trained transformer for computational chemistry. *Mach. Learn.: Sci. Technol.*, 3(1):015022, March 2022.
- [33] Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Min Wu, Tingjun Hou, and Mingli Song. Root-aligned SMILES: a tight representation for chemical reaction prediction. *Chem. Sci.*, 13(31):9023–9034, 2022.
- [34] Umit V. Ucak, Islambek Ashyrmamatov, Junsu Ko, and Juyong Lee. Retrosynthetic reaction pathway prediction through neural machine translation of atomic environments. *Nat Commun*, 13(1):1186, March 2022.
- [35] Yu Shee, Haote Li, Pengpeng Zhang, Andrea Nikolic, Sanil Sreekumar, Frédéric Buono, Jinhua Song, Timothy Newhouse, and Victor Batista. Site-specific template generative approach for retrosynthetic planning. *ChemRxiv*, 2024.
- [36] Marwin H. S. Segler, Mike Preuss, and Mark P. Waller. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature*, 555(7698):604–610, 2018.
- [37] Akihiro Kishimoto, Beat Buesser, Bei Chen, and Adi Botea. Depth-first proof-number search with heuristic edge cost and application to chemical synthesis planning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [38] Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro\*: Learning retrosynthetic planning with neural guided A\* search. In *The 37th International Conference on Machine Learning (ICML 2020)*, 2020.
- [39] Samuel Genheden and Esben Bjerrum. PaRoutes: towards a framework for benchmarking retrosynthesis route predictions. *Digital Discovery*, 1(4):527–539, 2022.
- [40] Amol Thakkar, Thierry Kogej, Jean-Louis Reymond, Ola Engkvist, and Esben Jannik Bjerrum. Datasets and their influence on the development of computer assisted synthesis planning tools in the pharmaceutical domain. *Chem. Sci.*, 11(1):154–168, 2020.
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [42] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023.
- [43] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [44] E. J. Corey and Xue-Min Cheng. *Logic of chemical synthesis*. Wiley, New York, new ed. edition, 1995.
- [45] Jiacheng Xiong, Wei Zhang, Zunyun Fu, Jiatao Huang, Xiangtai Kong, Yitian Wang, Zhaoping Xiong, and Mingyue Zheng. Improve retrosynthesis planning with a molecular editing language. preprint, Chemistry, 2023.

- [46] Masahiro Kajino, Atsushi Hasuoka, and Haruyuki Nishida. 1-heterocyclylsulfonyl, 2 aminomethyl, 5-(hetero-) aryl substituted 1-h-pyrrole derivatives as acid secretion inhibitors., WO2007026916A1, 2007.
- [47] Qian-Ying Yu, Huang Zeng, Kai Yao, Jian-Qi Li, and Yu Liu. Novel and practical synthesis of vonoprazan fumarate. *Synthetic Communications*, 47(12):1169–1174, 2017.
- [48] Jeffrey O. Saunders, Francesco G. Salituro, and Shunqi Yan. Therapeutic compounds and compositions, WO2011002817A1, 2011.
- [49] Jacob P. SIZEMORE, Liting GUO, Mahmoud Mirmehrabi, and Yeqing Su. Crystalline forms of n-(4-(4-(cyclopropylmethyl) piperazine-1-carbonyl)phenyl)quinoline-8-sulfonamide, WO2019104134A1, 2019.
- [50] Zhefeng Zhang, Aiqin Zhang, and Mouli Yang. Fused-ring compound having analgesic activity, and preparation method therefor and use thereof, WO2023160004A1, 2023.

## 8 Supplementary Information

Table 3: Top- $K$  accuracy on subset of test set- $n_1$  (random 500 routes).

Method	Top-1	Top-2	Top-3	Top-4	Top-5	Top-10
DirectMultiStep-SM-10M	0.402	0.476	0.522	0.550	0.564	0.582
DirectMultiStep-SM-60M	0.398	0.462	0.496	0.518	0.526	0.540
DirectMultiStep-noSM-10M	0.292	0.362	0.384	0.388	0.396	0.402
DirectMultiStep-noSM-60M	0.366	0.438	0.460	0.478	0.484	0.490

Table 4: Top- $K$  accuracy subset of test set- $n_5$  (random 500 routes).

Method	Top-1	Top-2	Top-3	Top-4	Top-5	Top-10
DirectMultiStep-SM-10M	0.374	0.442	0.474	0.484	0.498	0.516
DirectMultiStep-SM-60M	0.344	0.428	0.454	0.456	0.462	0.480
DirectMultiStep-noSM-10M	0.222	0.264	0.282	0.288	0.290	0.298
DirectMultiStep-noSM-60M	0.310	0.368	0.388	0.400	0.406	0.418

Table 3 and Table 4 present the Top- $K$  accuracy results for 500 random routes each from set- $n_1$  and set- $n_5$ . These results illustrate that a larger model is necessary to maintain high performance when starting material information is not provided, as shown by the DirectMultiStep-noSM-60M model's superior performance compared to DirectMultiStep-noSM-10M.

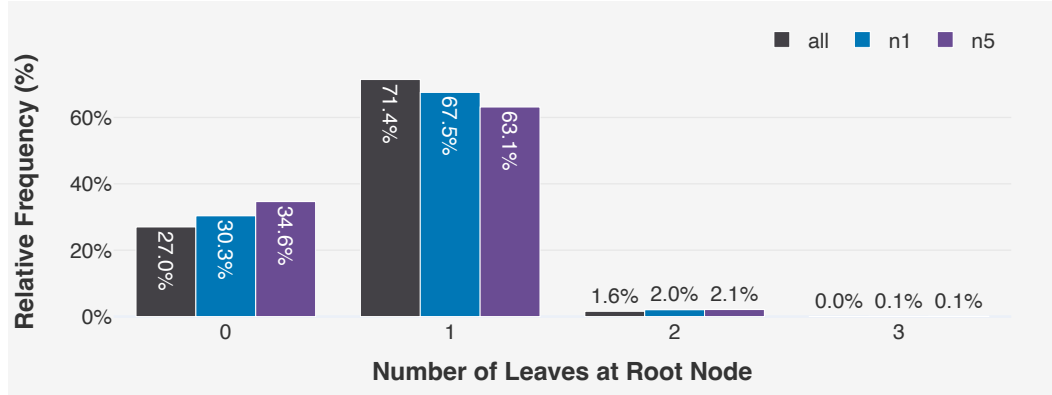


Figure 7: Distribution of the relative frequencies for the number of leaves (nodes with no children) at root node.