

RESEARCH

Open Access



Ualign: pushing the limit of template-free retrosynthesis prediction with unsupervised SMILES alignment

Kaipeng Zeng¹, Bo Yang², Xin Zhao¹, Yu Zhang¹, Fan Nie³, Xiaokang Yang¹, Yaohui Jin^{1*} and Yanyan Xu^{1*}

Abstract

Motivation Retrosynthesis planning poses a formidable challenge in the organic chemical industry, particularly in pharmaceuticals. Single-step retrosynthesis prediction, a crucial step in the planning process, has witnessed a surge in interest in recent years due to advancements in AI for science. Various deep learning-based methods have been proposed for this task in recent years, incorporating diverse levels of additional chemical knowledge dependency.

Results This paper introduces UAlign, a template-free graph-to-sequence pipeline for retrosynthesis prediction. By combining graph neural networks and Transformers, our method can more effectively leverage the inherent graph structure of molecules. Based on the fact that the majority of molecule structures remain unchanged during a chemical reaction, we propose a simple yet effective SMILES alignment technique to facilitate the reuse of unchanged structures for reactant generation. Extensive experiments show that our method substantially outperforms state-of-the-art template-free and semi-template-based approaches. Importantly, our template-free method achieves effectiveness comparable to, or even surpasses, established powerful template-based methods.

Scientific contribution We present a novel graph-to-sequence template-free retrosynthesis prediction pipeline that overcomes the limitations of Transformer-based methods in molecular representation learning and insufficient utilization of chemical information. We propose an unsupervised learning mechanism for establishing product-atom correspondence with reactant SMILES tokens, achieving even better results than supervised SMILES alignment methods. Extensive experiments demonstrate that UAlign significantly outperforms state-of-the-art template-free methods and rivals or surpasses template-based approaches, with up to 5% (top-5) and 5.4% (top-10) increased accuracy over the strongest baseline.

Keywords Template-Free Retrosynthesis Prediction, Deep Learning, Chemical Reactions, Single-step Retrosynthesis Prediction

*Correspondence:

Yaohui Jin

jinyh@sjtu.edu.cn

Yanyan Xu

yanyanxu@sjtu.edu.cn

Full list of author information is available at the end of the article



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

Introduction

Retrosynthesis prediction is a crucial task in organic chemistry, aiding in finding efficient synthetic pathways from target molecules to accessible starting materials. Despite significant advancements in chemical synthesis technology, it still remains a challenge in industries like pharmaceuticals. The extensive search space and the incomplete understanding of chemical reaction mechanisms make retrosynthesis prediction difficult, even for experienced chemists. To address this issue, computer-assisted synthetic planning (CASP) has gained increasing attention in recent years, starting from the seminal work by Corey. This paper focuses on single-step retrosynthesis prediction, which is the fundamental step in CASP. It aims to predict the reactants that can lead to a given product molecule through a single reaction step.

Various deep-learning-based single-step retrosynthesis prediction methods have been proposed in recent years. These methods can be broadly classified into three groups based on their dependency on additional chemical knowledge: template-based, semi-template-based and template-free methods. Template-based methods [5, 7, 9, 44] require an extra database of reaction templates. They frame the retrosynthesis prediction as a classification or retrieval problem for reaction templates suitable for the given product molecule to be synthesized. Among these solutions, Retrosim [7] utilizes molecular similarity to rank reaction templates; LocalRetro [5] and GLN [9] use graph neural networks to model the relationship between reaction templates and molecules to predict the most suitable reaction template; RetroKNN [44] further improves upon LocalRetro by addressing the issue of data imbalance using K-nearest neighbors (KNN). Template-based methods have strong interpretability and can accurately predict reactants. However, these methods are often unable to cover all cases and suffer from poor scalability due to limitations imposed by the template database.

To overcome the limitations faced by template-based methods, researchers have turned to generative models. Semi-template-based methods incorporate chemical knowledge into generative models with the help of chemical toolkits like RDKit [19], breaking free from the limitations imposed by reaction templates. The key idea of most semi-template-based methods [6, 32, 33, 41, 45] is to first convert the product into synthons based on reaction center identification and then complete the synthons into reactants. Graph neural networks are commonly used for synthon prediction, followed by leaving group attachment [6, 33], conditional graph generation [32], or SMILES (Simplified Molecular Input Line Entry System) generation [45] for reactant completion. Apart from all above, RetroPrime [41] utilizes two independent

Transformers to accomplish synthon prediction and reactant generation as separate tasks.

Semi-template-based methods to a certain extent are more in line with chemical intuition. However, these methods increase the complexity of inference and training as they break down retrosynthesis into two subtasks. Failures in synthon prediction directly affect subsequent reactant completion and overall performance. Besides, methods based on leaving group necessitates an extra leaving group database. This requirement, akin to template-based approaches, imposes limitations on the model's scalability.

As generative models, Template-free methods opt to generate reactants directly from the given products. In comparison to generating graph structures, SMILES provides a way to represent molecules as strings. Taking advantage of this, most template-free methods [18, 31, 34, 40, 49] use Transformer models to translate between product SMILES and reactants SMILES. In particular, Graph2SMILES [35] replaces the Transformer encoder with a graph neural network, resulting in a permutation-invariant pipeline. There are also methods [27, 47] formulates the generation of reactants as a series of graph generation or editing operation and solve it autoregressively. Existing template-free methods generally follows an auto-regressive generation strategy and use beam search for the generation process. Consequently, preserving a level of diversity in the resultant outputs has emerged as a critical consideration for template-free methods [39]. Due to the use of SMILES as input and output, most of template-free methods often overlook the rich topological and chemical bond information present in molecular graphs. Moreover, as reactants molecules need to be generated from scratch, template-free methods frequently suffer from validity issues and fail to leverage an important property of retrosynthesis prediction, i.e., the presence of many common substructures between products and reactants.

In this paper, we focus on the template-free generative approach for retrosynthesis prediction. Existing sequence-to-sequence methods have limitations in extracting robust molecular representations. They overlook the abundance of topological information and chemical bonds, and lack the ability to utilize atom descriptors as rich as those in graph-based methods. Furthermore, template-free methods overlook the fact that the molecular graph topology remains largely unaltered from reactants to products during chemical reactions, as they generate reactants from scratch. While there are methods that attempt to solve this problem using supervised SMILES alignment, they require complex data annotation and impact model training. Given these limitations, the following question naturally arises:

Can we effectively leverage the structural information of product molecules using a much simpler approach?

To address these issues and further enhance template-free methods, we propose a novel graph-to-sequence pipeline called UAlign. Our approach employs a specifically designed graph neural network as an encoder, incorporating information from chemical bonds during message passing to create more powerful embeddings for the decoder. We introduce an unsupervised SMILES alignment mechanism that establishes associations between product atoms and reactant SMILES tokens which reduces the complexity of SMILES generation and enables the model to focus on learning chemical knowledge. Our model outperforms existing template-free

methods by a large margin and demonstrates comparable performance against template-based methods.

Methods

We introduce UAlign, a novel single-step retrosynthesis prediction model based on an encoder-decoder architecture, as demonstrated in Fig. 1. It's a fully template-free method without any molecule editing operation using RDKit [19]. We propose a specially designed variant of Graph Attention Network, which incorporates the information of chemical bonds to enhance the capability of capturing the structural characteristics of molecules.

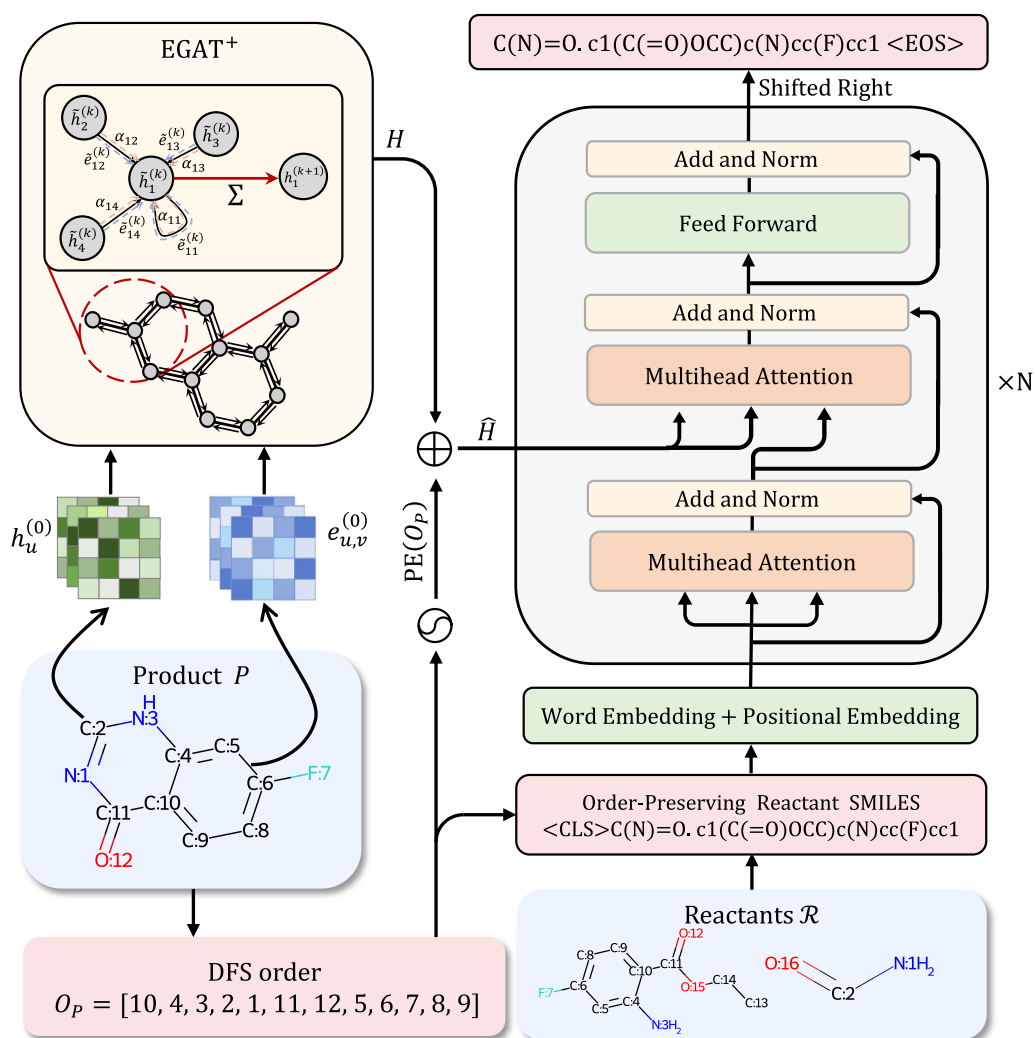


Fig. 1 Overview of UAlign: Given a product molecule graph P and one of its DFS order O_p , the graph is first fed into the graph neural network called EGAT+ to obtain node features H . Then the positional encoding is added to H according to the given DFS order O_p to generate the order-aware node features \hat{H} . Finally the decoder takes \hat{H} as input and generate the SMILES of reactants auto-regressively

Preliminary

A molecule can be represented as a graph, denoted by $G = (V, E)$, where V represents the atoms and E represents the chemical bonds. The SMILES representation of a molecule can be obtained by performing a depth-first search (DFS) starting from any arbitrary atom in the molecule graph. Given a molecule graph $G = (V, E)$, we can generate multiple DFS orders and each DFS order corresponds to a SMILES representation of the graph. Denoted the set of all possible DFS orders as $\mathcal{D}(V) \subseteq \mathcal{P}(V)$, $\mathcal{P}(V)$ represents all permutations of the set of atoms V . For each DFS order $O \in \mathcal{D}(V)$, we denoted its corresponding SMILES as $Smiles(G, O)$, which lists all atoms in the order dictated by O . To facilitate our subsequent elaboration, we refer to the position of an atom a in the order O as its rank, denoted as $rank(a, O)$. The atom with the minimal rank given order O is then defined as the root atom, denoted as $root(G, O)$.

Model architecture

In this section, we provide an overview of the model's architectural design and the rationale behind it. Our model adopts an encoder-decoder framework, as depicted in Fig. 1. The encoder is tasked with extracting molecular representations from the input products and supplying these as inputs to the decoder, which then generates a combination of reactants. We utilize graph neural networks for our encoder, which encodes the nodes through an iterative message-passing mechanism to derive node features. During each round of message-passing, the network collects and aggregates information from a node's neighbors, thereby updating the node features. This design effectively integrates the topological information of the graph structure into the node features, naturally adapting to the task of molecular representation learning. The superiority of graph neural networks in this domain has been substantiated by a plethora of studies [15–17].

Moving on to the decoder, it is constructed based on the transformer decoder architecture [37]. Generating a graph poses a unique challenge due to the lack of inherent order among the graph's nodes and the necessity to predict an adjacency matrix that is quadratic in relation to the number of nodes. However, the adoption of SMILES [42] circumvents this complexity by converting the molecular generation problem into a more tractable text sequence generation task. This transformation is crucial because it ensures that the length of the output sequence is linearly proportional to the number of atoms involved. And the framework of text sequence generation has been extensively applied in other domains, such as natural language processing, providing a robust

foundation upon which our molecular generation model is built [12, 21, 22, 25]. The transformer decoder, equipped with a cross-attention mechanism, is adept at sequential generation conditioned on a given input, making it an ideal choice for our model's decoder. The subsequent sections will delve into the intricate design of the encoder, decoder, and other components of our model.

EGAT⁺

Chemical bonds play a significant role in determining the properties of molecules and contain valuable information. Previous studies [13, 26, 46] have demonstrated that incorporating edge information into graph neural networks can greatly enhance their ability to represent molecular structures. To fully leverage the information brought by chemical bonds, we propose a modified version of the Graph Attention Network (GAT) [38] called EGAT⁺.

Our proposed model explicitly incorporates edge features, which represent the information derived from chemical bonds, into the message passing process. During each iteration of message passing, the EGAT⁺ applies self-attention to each node and its one-hop neighbors to calculate attention coefficient according to both node features and edge features. It then aggregates the both node and edge features of these neighbors, considering the attention coefficients, to update the node features. Denote the node feature of atom u as $h_u^{(k)}$ and the edge feature between atom u and v as $e_{u,v}^{(k)}$ at k -th iteration of message passing. In math, the message passing mechanism can be written as

$$\begin{aligned} \tilde{e}_{u,v}^{(k)} &= \text{FFN}_e^{(k)}(e_{u,v}^{(k)}), \\ \tilde{h}_u^{(k)} &= \text{FFN}_n^{(k)}(h_u^{(k)}), \\ c_{u,v} &= \mathbf{a}^T [\tilde{h}_u^{(k)} \| \tilde{h}_v^{(k)} \| \tilde{e}_{u,v}^{(k)}], \\ \alpha_{u,v} &= \frac{\exp(\text{LeakyReLU}(c_{u,v}))}{\sum_{v' \in \mathcal{N}(u) \cup \{u\}} \exp(\text{LeakyReLU}(c_{u,v'}))}, \quad (1) \\ h_u^{(k+1)} &= \sum_{v \in \mathcal{N}(u) \cup \{u\}} \alpha_{u,v} (\tilde{h}_u^{(k)} + \tilde{e}_{u,v}^{(k)}), \\ e_{u,v}^{(k+1)} &= \text{FFN}_m^{(k)}([h_u^{(k+1)} \| h_v^{(k+1)} \| e_{u,v}^{(k)}]), u \neq v, \end{aligned}$$

where $\text{FFN}_m^{(k)}$, $\text{FFN}_e^{(k)}$ and $\text{FFN}_n^{(k)}$ are three different feed forward networks, \mathbf{a} is a learnable parameter, $\mathcal{N}(u)$ denotes the one-hop neighbors of node u and $\|$ denotes the concatenation operation. Since there are no chemical bonds with the same beginning and ending atoms, the $e_{u,u}^{(k)}$ is also set as a learnable parameter shared among all atoms. The residual connection and layer normalization [2] are applied to prevent over-smoothing while enlarging the receptive field of the model [43].

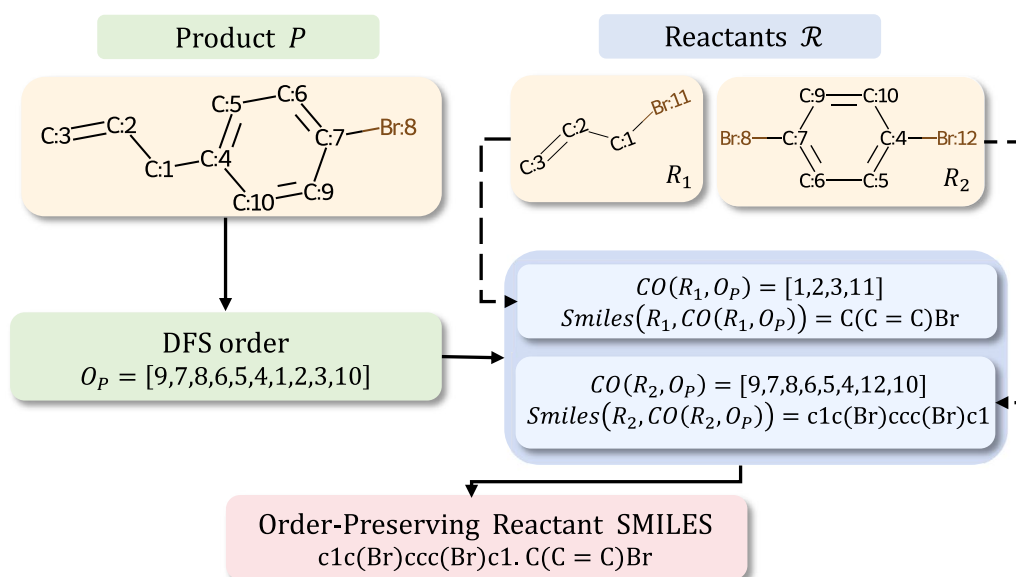


Fig. 2 An example of the process to generate order-preserving reactants SMILES. The atom mapping numbers shown on the figure are included only for clearer explanation and will be removed in our implementation to prevent any label leakage

The initial node features $h_u^{(0)}$ and edge features $e_{u,v}^{(0)}$ are determined via several chemical property descriptors, whose details are shown in Supplementary Sec. 6. After K iterations of message passing, we can obtain the encoded features $h_u^{(K)}$ of all atoms and make up the output $H \in \mathbb{R}^{V_P \times d}$ of the encoder, where d denotes the embedding size.

SMILES alignment

For single-step retrosynthesis prediction, a significant proportion of structures are shared between product molecules and reactant molecules [40, 50]. However, SMILES-based methods often have to generate the reactant SMILES from scratch, even if most of the structures of reactants are the same as those of the products. This results in the underutilization of input information and becomes the bottleneck of template-free retrosynthesis prediction methods. There are methods [31, 40] addressing this issue through supervised SMILES alignment, which involves adding supervised information to establish the correspondence between input and output tokens through cross-attention over the input and predicted tokens. This supervised training approach not only requires complex data annotation algorithms but also limits the diversity of the model's attention map, thereby further affecting the model's performance. To address the above-mentioned issues, we propose the unsupervised SMILES alignment method as follows.

Assuming we can identify the location of each product atom in the reactants' SMILES and provide it to the model, a natural correspondence can be established between the input and output atoms. However, during the inference process, revealing this information would lead to label leakage, which is not permitted. Therefore, we propose the following modification: when providing an order of product atoms, we expect the model to generate atom tokens in the reactants' SMILES in this given order as closely as possible. By doing so, we can establish a correspondence between the product atoms and the reactants' SMILES tokens using unsupervised methods without leaking any labels. We refer to this type of reactants' SMILES, which aims to preserve the given order of atom tokens as much as possible, as order-preserving reactant SMILES. Note that SMILES represents atoms in a molecule according to a certain DFS order, the provided order should also be a DFS order for the product molecule.

The generation of order-preserving reactant SMILES will be introduced as follows. Given the product molecule $P = (V_P, E_P)$ with a DFS order $O_P \in \mathcal{D}(V_P)$ and the corresponding set of reactant molecules $\mathcal{R} = \{R_1, R_2, \dots, R_I\}$, for each reactant $R = (V_R, E_R) \in \mathcal{R}$, we can find a depth-first order $O_R \in \mathcal{D}(V_R)$ that has a nearly consistent atomic appearance sequence with O_P as the product and reactants have similar structures. For convenience, we name such a order as O_P -corresponding order of R and denote it as $CO(R, O_P)$. Mathematically, it's defined as

$$CO(R, O_p) = \arg \min_{o \in \mathcal{D}(R)} \sum_{i \in O_p \cap o} \sum_{j \in O_p \cap o} inv(i, j, O_p, o), \quad (2)$$

where the value of $inv(i, j, O_p, o)$ equals 1 if and only if $rank(i, O_p) < rank(j, O_p)$ and $rank(i, o) > rank(j, o)$, and equals 0 otherwise. We sort the reactants \mathcal{R} according to $rank(root(R, CO(R, O_p)), O_p)$ in ascending order. Then we generate SMILES for each reactant molecule using its O_p -corresponding order and join them together using “.” to obtain order-preserving reactant SMILES.

For further discussion, we denote the order-preserving reactant SMILES given the reactant molecules \mathcal{R} and a DFS order O of product as $OPSmiles(\mathcal{R}, O)$. An example of the process to generate order-preserving reactants SMILES is shown in Fig. 2. The detailed implementations are presented in Supplementary Sec. 5.1.

Decoder

The decoder takes the node features $H \in \mathbb{R}^{V_p \times d}$ that are generated from the encoder, as well as the given DFS

order O_p for the product molecule graph as input. We use the vanilla Transformer decoder [37] as our decoder. As mentioned in "SMILES alignment" Section, the order information of product atoms are required for SMILES alignment. However, the Transformer decoder is permutation-invariant to memory [20, 35], meaning it is not sensitive to the order of the features from encoder. This implies that directly performing cross-attention over H may not effectively capture the relationship between product atoms and reactant SMILES tokens. To address this problem, we introduce position encoding to the node features based on the rank of each atom in the given DFS order O_p to generate order-aware node features \hat{H} . Then given an input embedding sequence $Z \in \mathbb{R}^{m \times d}$, the Transformer decoder layer utilizes the order-aware node features as keys and values in all the cross-attention layers. This process ultimately generates the decoded embeddings $\hat{Z} \in \mathbb{R}^{m \times d}$. These embeddings are then fed into feed-forward layers $FFN_1: \mathbb{R}^d \rightarrow \mathbb{R}^T$ to predict

Table 1 Top- k accuracy for retrosynthesis prediction on USPTO-50K

Model	Top- k accuracy (%)							
	Reaction class known				Reaction class unknown			
	1	3	5	10	1	3	5	10
Template-Based								
RetroSim [7]	52.9	73.8	81.2	88.1	37.3	54.7	63.3	74.1
NeuralSym [30]	55.3	76.0	81.4	85.1	44.4	65.3	72.4	78.9
GLN [9]	64.2	79.1	85.2	90.0	52.5	69.0	75.6	83.7
LocalRetro [5]	63.9	86.8	92.4	96.3	53.4	77.5	85.9	92.4
RetroKNN [44]	66.7	88.2	93.6	96.6	57.2	78.9	86.4	92.7
Semi-Template-Based								
RetroXpert* [45]	62.1	75.8	78.5	80.9	50.4	61.1	62.3	63.4
G2G [32]	61.0	81.3	86.0	88.7	48.9	67.6	72.5	75.5
GraphRetro [33]	63.9	81.5	85.2	88.1	53.7	68.3	72.2	75.5
RetroPrime* [41]	64.8	81.6	85.0	86.9	51.4	70.8	74.0	76.1
Template-Free								
Transformer [37]	57.1	71.5	75.0	77.7	42.4	58.6	63.8	67.7
SCROP [49]	59.0	74.8	78.1	81.1	43.7	60.0	65.2	68.7
Liu's Seq2seq [24]								
Tied Transformer [18]	–	–	–	–	37.4	52.4	57.0	61.7
Aug. Transformer* [34]	–	–	–	–	47.1	67.1	73.1	76.3
MEGAN [27]	60.7	82.0	87.5	91.6	48.1	70.7	78.4	86.1
GTA* [31]	–	–	–	–	51.1	67.6	74.8	81.6
Graph2SMILES [35]	–	–	–	–	52.9	66.5	70.0	72.9
Retroformer* [40]	64.0	82.5	86.7	90.2	53.2	71.1	76.6	82.1
RetroBridge [14]	–	–	–	–	50.8	74.1	80.6	85.6
Ours*- Average	66.4	86.7	91.5	95.0	53.5	77.3	84.6	90.5
Ours*- Standard Deviation	0.1	0.2	0.2	0.1	0.2	0.3	0.3	0.2

For comparison purpose, the Aug. Transformer is evaluated without the test augmentation. Best performance of each model type is in bold. The average and standard deviation of our method over five different random seeds are reported

* indicates the model with SMILES augmentation.

the tokens \hat{T} that should be generated. In summary, the decoder can be mathematically expressed as

$$\begin{aligned}\hat{H} &= H + \text{PE}(O_P), \\ \hat{Z} &= \text{TransformerDecoder}(Z, \mathbf{A}), \\ \hat{T} &= \text{FFN}_1(\hat{Z}).\end{aligned}\quad (3)$$

Two-stage training

There is a significant distribution shift between graphs and SMILES representations. Moreover, our model is specifically designed to generate non-canonical SMILES, which may contain more complex patterns compared to canonical SMILES. To achieve this goal, we propose a two-stage training strategy in this paper. The first stage aims to align the distributions between two distinct modalities: SMILES and molecular graphs, while enabling the model to learn the patterns of non-canonical SMILES. Given a molecule graph M and one of its possible DFS orders O_M , the training task is to translate the graph into the corresponding SMILES representation based on the given order O_M . In detail, this is reached by training the model to generate $\text{Smiles}(M, O_M)$ given molecule M and DFS order O_M .

Once the first stage training converges, we proceed to the second stage, which focuses on retrosynthesis prediction. In this stage, the model is trained using the order-preserving reactant SMILES as targets. Given a product molecule graph P , a possible DFS order O_P , and a set of reactants \mathcal{R} , the model is expected to generate $\text{OPSmiles}(\mathcal{R}, O_P)$.

Data augmentation

Different from those Transformer-based methods [31, 34, 40] taking SMILES as input and canonical SMILES as target, our method takes a graph as input and is trained with non-canonical SMILES. That means the previous SMILES augmentation tricks are not suitable for us. Similar to [40], we choose to augment the training data on-the-fly.

For the first stage, at each iteration, for each molecule $M = (V_M, E_M)$, we have a 50% chance of using a random DFS order O_M as the input for the model, and using the corresponding $\text{Smiles}(M, O_M)$ as the training target. For the other 50%, we randomly select another molecule $M' = (V_{M'}, E_{M'})$ from the dataset to form a new molecular graph $\tilde{M} = (V_M \cup V_{M'}, E_M \cup E_{M'})$, and find the DFS order $O_{\tilde{M}}$ that can generate canonical SMILES for \tilde{M} . \tilde{M} and $O_{\tilde{M}}$ are then fed into the model and the target is set as the canonical SMILES of \tilde{M} . Such an augmentation method enables the model to output the atom tokens according to the given DFS order and be aware of different components within a graph.

For the second stage, at each iteration, for each product molecule $P = (V_P, E_P)$, we have a 50% probability of using a random DFS order as input, and for the remaining part, we use the DFS order capable of producing canonical SMILES for product as input. The target used for training is the order-preserving reactant SMILES generated based on the input DFS order. This data augmentation method allows the model to focus more on the DFS order for canonical product SMILES while also noticing the correspondence between product atoms and the output SMILES tokens.

Loss function

Both of the two stages of training can be considered as a kind of translation between graphs and SMILES, thus we use the loss widely used for auto-regressive language generation models for training. Denote the training target as $T = \{t_1, t_2, \dots, t_n\}$ and the output of the model $\tilde{T} = \{\hat{t}_1, \hat{t}_2, \dots, \hat{t}_n\}$, the loss can be written as

$$\mathcal{L} = \sum_{i=1}^n l_{cls}(\hat{t}_i, t_i), \quad (4)$$

where $l_{cls}(\cdot)$ is the classification loss.

Moreover, we have introduced numerous definitions to establish a foundation for understanding the model's overall training process. To facilitate comprehension, we have encapsulated the entire training procedure in Algorithm 1.

Algorithm 1 The training procedure.

Input: The set of all the molecules in train set $\mathcal{D}_m = \{G_i\}_{i=1}^{N_{mol}}$, all the reaction of train set $\mathcal{D}_r = \{(P_i, \mathcal{R}_i)\}_{i=1}^{N_{train}}$, number of training epochs of the first stage E_1 , number of training epochs of the second stage E_2 , batch size B .

Output: Trained parameter θ .

for $i \leftarrow 1$ **to** $E_1 \times \frac{N_{mol}}{B}$ **do** ▷ The first training stage
 Sample a batch $\{G_{b_1}, G_{b_2}, \dots, G_{b_B}\}$ with size B from \mathcal{D}_m .
 Initialize input graphs \mathcal{G}_I as empty list.
 Initialize input DFS orders \mathcal{D}_I as empty list.
 Initialize targets \mathcal{T} as empty list.
for $j \leftarrow 1$ **to** B **do**
 $prob \leftarrow Uniform(0, 1)$.
 if $prob \leq 0.5$ **then**
 Generate random DFS order D for G_{b_j} .
 Generate corresponding SMILES of DFS order D as $Smiles(G_{b_j}, D)$.
 Add D to \mathcal{D}_I , add G_{b_j} to \mathcal{G}_I , add $Smiles(G_{b_j}, D)$ to \mathcal{T} .
 else
 Random Selecte $G \in \mathcal{D}_{mol}$.
 Merge G_{b_j} and G as a new molecule graph G_m .
 Generate DFS order D for canonical SMILES of G_m .
 Generate canonical SMILES of G_m as $Smiles(G_m, D)$.
 Add D to \mathcal{D}_I , add G_m to \mathcal{G}_I , add $Smiles(G_m, D)$ to \mathcal{T} .
 end if
end for
 Generate model output $\hat{\mathcal{T}}$ with input graphs \mathcal{G}_I and DFS orders \mathcal{D}_I .
 Backpropagate $\mathcal{L}(\hat{\mathcal{T}}, \mathcal{T})$ and optimize parameters θ .
end for

for $i \leftarrow 1$ **to** $E_2 \times \frac{N_{train}}{B}$ **do** ▷ The second training stage
 Sample a batch $\mathcal{B} = \{(P_{b_1}, \mathcal{R}_{b_1}), (P_{b_2}, \mathcal{R}_{b_2}), \dots, (P_{b_B}, \mathcal{R}_{b_B})\}$ with size B from \mathcal{D}_r .
 Initialize input DFS orders \mathcal{D}_I as empty list.
 Initialize targets \mathcal{T} as empty list.
for $j \leftarrow 1$ **to** B **do**
 $prob \leftarrow Uniform(0, 1)$.
 if $prob \leq 0.5$ **then**
 Generate DFS order D for canonical SMILES of P_{b_j} .
 else
 Generate random DFS order D for P_{b_j} .
 end if
 Generate order-preserving reactant SMILES of \mathcal{R}_{b_j} as $OPSmiles(\mathcal{R}_{b_j}, D)$.
 Add D into \mathcal{D}_I , add $OPSmiles(\mathcal{R}_{b_j}, D)$ into \mathcal{T} .
end for
 Generate model output $\hat{\mathcal{T}}$ with input graphs \mathcal{B} and DFS orders \mathcal{D}_I .
 Backpropagate $\mathcal{L}(\hat{\mathcal{T}}, \mathcal{T})$ and optimize parameters θ .
end for

Table 2 Top-*k* accuracy for retrosynthesis prediction on USPTO-MIT

Model Type	Model	Top- <i>k</i> accuracy (%)			
		1	3	5	10
template-based	NeuralSym [30]	47.8	67.6	74.1	80.2
	LocalRetro [5]	54.1	73.7	79.4	84.4
template-free	Liu's Seq2Seq [24]	46.9	61.6	66.3	70.8
	AutoSynRoute [23]	54.1	71.8	76.9	81.8
	RetroTRAE [36]	58.3	-	-	-
	Ours	59.9	76.9	82.0	86.4

Best performance of each model type is in bold

Table 3 Top-*k* accuracy for retrosynthesis prediction on USPTO-FULL

Model Type	Model	Top- <i>k</i> accuracy (%)			
		1	3	5	10
Template-based	RetroSim [7]	32.8	-	-	56.1
	NeuralSym [30]	35.8	-	-	60.8
	GLN [9]	39.3	-	-	63.7
	LocalRetro [5]	39.1	53.3	58.4	63.7
Semi-template-based	RetroPrime* [41]	44.1	59.1	62.8	68.5
Template-free	MEGAN [27]	33.6	-	-	63.9
	Aug. Transformer* [34]	46.2	-	-	73.3
	Graph2SMILES [35]	45.7	-	-	63.4
	GTA* [31]	46.6	-	-	70.4
	Ours*	50.4	66.1	71.3	76.2

Best performance of each model type is in bold

* indicates the model with SMILES augmentation.

Results and discussion

In this section, we conduct extensive experiments to make a comprehensive evaluation of our proposed UAlign.

Evaluation protocol

Benchmark Datasets. We adopt three datasets for evaluation: (1) USPTO-50K consists of 50,016 atom-mapped reactions grouped into 10 different classes; (2) USPTO-FULL comprises 1,013,118 atom-mapped reactions without any reaction class information. (3) **USPTO-MIT** consists of 479,035 atom-mapped reactions without any reaction class information. To ensure a fair comparison, we adopt the same training/validation/test splits as those in a previous study [9] for USPTO-50K and USPTO-FULL datasets. The training/validation/test splits is aligned with the previous study [16]. The detailed data processing procedure and the statistical information of the processed dataset are presented in Supplementary Sec. 2.

Metrics. We utilize the following three evaluation metrics for evaluation: top-*k* accuracy, top-*k* SMILES validity, top-*k* round-trip accuracy and Computational Cost. The detailed definitions for the first three metrics are provided in Supplementary Sec. 3.

Performance comparison

Top-*k* Accuracy. We compare our model with existing single-step retrosynthesis prediction in terms of top-*k* accuracy on all the datasets. The results are summarized in Tables 1, 2 and 3. On the USPTO-50K dataset, our model achieves an average top-3 accuracy of 77.3%, average top-5 accuracy of 84.6% and average top-10 accuracy of 90.5% under the reaction class unknown setting, surpassing the SOTA template-free method by 3.2%, 4.0% and 4.9% respectively. And with reaction class given on USPTO-50K dataset, our model achieves an average top-3 accuracy of 86.7%, average top-5 accuracy of 91.5% and average top-10 accuracy of 95.0%, which exceeds the SOTA template-free method by 4.2%, 4.0% and 4.8% respectively. It is worth noting that, even when taking the standard deviation into consideration, the lower-bound performance of our model still surpasses all template-free methods in terms of all metrics. Moreover, our model outperforms all the semi-template-based methods with a noticeable margin. It's also encouraging to see that our method, as a template-free method, achieves competitive or even superior performance against the powerful template-based methods such as LocalRetro under both settings of USPTO-50K dataset. On UPSTO-MIT dataset, our model achieves the top-1 accuracy of 59.9% and top-10 accuracy of 86.4%, which even outperforms the existing template-based SOTA method LocalRetro significantly. Additionally, our models achieved a top-1 accuracy of 50.4% on the USPTO-FULL dataset, which exceed that of the current SOTA model GTA by 3.8%. These findings sufficiently demonstrate the effectiveness of our method. The contribution of each proposed module will be further validated in "Ablation Study" section.

Table 4 Top-*k* SMILES validity for retrosynthesis prediction on USPTO-50K with reaction class unknown

Model	Top- <i>k</i> validity (%)			
	1	3	5	10
Transformer [37]	97.2	87.9	82.4	73.1
Graph2SMILES [35]	99.4	90.9	84.9	74.9
RetroPrime [41]	98.9	98.2	97.1	92.5
Retroformer [40]	99.2	98.5	97.4	96.7
Ours	99.8	99.7	99.3	98.2

Table 5 Top-*k* round-trip accuracy for retrosynthesis prediction on USPTO-50K with reaction class unknown

Model	Top- <i>k</i> round-trip acc. (%)			
	1	3	5	10
GraphRetro [33]	80.5	73.3	68.3	59.3
Transformer [37]	71.9	54.7	46.2	35.6
Graph2SMILES [35]	76.7	56.0	46.4	34.9
RetroPrime [41]	79.6	59.6	50.3	40.4
Retroformer [40]	78.9	72.0	67.1	57.2
Ours	80.9	74.0	69.0	60.2

Table 6 Average inference time per sample on USPTO-50K dataset

Model	Avg. Inference Time (s)
GraphRetro [33]	0.344
Graph2SMILES [35]	0.037
RetroPrime [41]	1.359
Retroformer [40]	0.803
Ours	0.481

It is noteworthy that while template-based approaches have achieved remarkable performance on the USPTO-50K dataset, their reliance on external template libraries has emerged as a constraint as datasets grow in scale and complexity. This dependency leads to a substantial degradation in model performance. In contrast, template-free methods have demonstrated superior versatility and adaptability, qualities that render them especially appropriate for managing large-scale and intricate datasets.

Top-*k* SMILES Validity. We use vanilla Transformer, RetroPrime, Retroformer and Graph2SMILES as robust baselines to compare the validity of SMILES in our study. SMILES generation models for retrosynthesis tasks often encounter challenges with maintaining SMILES

validity. We do not take the methods based on templates or molecule editing as baseline here because the validity of generated SMILES can be guaranteed by the templates or chemical toolkits. Unlike graph-based models, SMILES-based methods need to ensure that the generated content adheres to the parsing rules of SMILES, without leveraging chemical tools such as RDKit. Consequently, SMILES-based approaches are more susceptible to generating invalid SMILES compared to graph-based models. As shown in Table 4, our model demonstrates superior top-1 and top-5 molecule validity compared to other models, even without employing canonical SMILES as our training objective. This improvement could be attributed to the proposed two-stage training strategy and data augmentation, which assist the model in capturing various SMILES patterns effectively.

Top-*k* Round-Trip Accuracy. To assess the accuracy of our predicted synthesis plans, we utilize the Molecule Transformer [29] as the benchmark reaction prediction model and calculate the top-*k* round-trip accuracy. We take RetroPrime, Retroformer and Graph2SMILES as our strong SMILES-based baselines. We also use take graph-based method GraphRetro into comparison. The results are presented in Table 5. The results clearly indicate that our model outperforms all SMILES-based baselines by a considerable extent and even exceeds the well-established graph-based method, GraphRetro. This underscores the efficacy of our unsupervised SMILES alignment mechanism, which enables the model to efficiently leverage substructures from product molecules to construct reactants. This mechanism allows the model to focus more intently on learning reaction mechanisms, thereby yielding more plausible predictive outcomes. In summary, our model has exhibited a robust capacity for generating coherent and efficacious synthesis pathways, specifically tailored for advanced downstream applications such as multi-step retrosynthesis planning.

Computational Cost. The computational cost is a critical metric for single-step prediction models, particularly when these models are intended to be integrated

Table 7 Effects of different modules on retrosynthesis performance in reaction class unknown setting of USPTO-50K dataset

Method	Top- <i>k</i> acc. (%)				Top- <i>k</i> round-trip acc. (%)				Top- <i>k</i> validity (%)			
	1	3	5	10	1	3	5	10	1	3	5	10
UAlign (Full Version)	53.6	77.6	84.6	90.3	80.9	74.0	69.0	60.2	99.8	99.7	99.3	98.2
– Two-Stage Training	53.1	75.6	82.8	88.8	80.4	72.3	66.9	57.6	99.2	97.7	96.1	92.2
– Data Augmentation	48.5	71.8	79.2	85.9	78.2	70.5	65.1	56.0	99.1	98.4	97.7	95.6
– SMILES Alignment	45.5	65.2	71.4	77.5	70.3	57.9	51.0	40.1	97.8	95.3	93.6	89.2
Transformer [37]	42.4	58.6	63.8	67.7	71.9	54.7	46.2	35.6	97.2	87.9	82.4	73.1

Best performance is in bold

with other searching algorithms for multi-step retrosynthesis planning and are expected to be invoked repeatedly. In our comparative analysis, we included SMILES-based baselines including RetroPrime, Retroformer, Graph2SMILES, and the graph-based baseline, GraphRetro. We performed inference on the test set of the USPTO-50K dataset using a single NVIDIA RTX 3090 graphics card. The average inference time per sample is detailed in Table 6.

As it is shown in the table, our model exhibits the second-fastest inference speed among the SMILES-based methods and shows a negligible difference in inference time when compared to the graph-based baseline, GraphRetro. The superior inference speed of Graph2SMILES is attributed to its use of an RNN as the decoder, which has a computational complexity linearly related to the length of the output sequence. In contrast, the other SMILES-based methods are constructed with a transformer decoder, resulting in a computational complexity that is quadratically related to the output sequence length. These results underscore the capability of our method to rapidly infer results based on input, positioning it favorably for integration with searching algorithms that require extensive exploration and trial-and-error in the construction of multi-step retrosynthesis planning systems.

It is noteworthy that our method does not implement batch-wise parallelism and only utilized 2GB of GPU memory during inference. There is significant room for

optimization in our code, which could enable parallel inference for multiple samples, thereby achieving greater hardware utilization efficiency and faster average inference speeds.

Ablation study

We investigate the effects of different components in our proposed pipelines. The result is summarized in Table 7.

Two-Stage Training. We eliminate the initial training phase and directly train the model for the retrosynthesis prediction task. As indicated in Table 7, the two-stage training strategy has consistently led to enhancements in all evaluated metrics. This observation implies that the two-stage training strategy effectively enables the model to adeptly learn the intricacies of molecular SMILES representations, thereby yielding higher quality and more plausible retrosynthetic analysis outcomes.

Data Augmentation. We remove the data augmentation during the second training stage, which means training solely using the DFS order that can generate canonical SMILES. Table 7 demonstrates a significant decline in model performance across all metrics. This clearly demonstrates that our data augmentation significantly improves the model's performance.

SMILES Alignment. In the training process, we remove all operations related to SMILES alignment. This includes the removal of the position encoding in Eq. 3, where the features H directly served as the input memory for




Fig. 3 Visualization of cross-attention over order-aware node features and the predicted tokens. The number on the y-axis is the map number of atoms in the product. The reactant atoms that not appear in product is colored red in the x-axis. o represents the end token

the Transformer decoder. Since we eliminate the input related to the DFS order, the model was no longer trained using order-preserving reactants SMILES as the target but instead switched to canonical SMILES for product. Additionally, in this set of experiments, we remove the first training stage, which aligns the graph and SMILES modalities as the model architecture changes. The results are reported in Table 7, and they show a significant decline in performance compared to our full version, indicating that the proposed SMILES alignment algorithm is crucial for achieving excellent performance.

It is worth noting that even without data augmentation, two-stage training and SMILES alignment, our model still outperforms the vanilla Transformer by a large margin in terms of all metrics reported in the last line of Table 7. This indicates that graph-based molecular representation learning still has advantages over SMILES-based approaches, and our proposed EGAT⁺ can extract effective molecular representations for downstream usage.

Case study (visualization of cross-attention mechanism in transformer with UAlign)

We randomly select a case from the dataset and showcase the cross-attention map in Fig. 3. The cross-attention map indicates the correlation between reactant tokens and nodes in the input product graph. This map is obtained by averaging the attention coefficient from each attention head. From the figure, it is evident that the predicted tokens successfully locate their corresponding atoms in the product, which contributes to the accurate prediction. The SMILES alignment can also be observed to assist the model in correctly identifying the reaction center. In accordance with the figure, the bond between atom C:11 and N:9 breaks during the transformation into reactants. Our model effectively notices this occurrence and focus the attention of token \hat{t}_7 on the reaction centers C:11 and N:9. This strategic focusing successfully guides the completion of the reactants, ensuring that the leaving group is correctly attached to the appropriate atoms. Additionally, we note that the attention coefficient at token \hat{t}_{14} is concentrated on atoms C:1 and N:9, which are the first atoms of each reactant molecule according to the given DFS order. This further indicates that our model is able to correctly identify the sites where the reaction occurs and accurately cleave the chemical bonds. Moreover, the attention of newly generated structures (i.e., tokens \hat{t}_7 to \hat{t}_{13}) is directed towards atoms C:1, C:11, and O:2, which correspond to the specific synthon they will attach to. This demonstrates that our model is able to generate appropriate functional groups based on the molecular structure information to form the reactants. All the aforementioned results illustrate

that our proposed SMILES alignment method assists the model in comprehending molecular structural information and helps it to focus on learning chemical rules.

To further investigate the impact of the proposed SMILES alignment mechanism on model training, we visualize the cross-attention coefficients of different Transformer decoder layers. The visualization is provided in Supplementary Fig. 1 of Supplementary Information. From Supplementary Fig. 1, we can observe significant variations in the cross-attention across different layers. Additionally, the establishment of correspondence does not occur exclusively at certain layers, such as the first or last layer. This suggests that directly imposing supervised signals on the cross attention coefficient [31, 40] for SMILES alignment is not a wise approach, whether applied to all layers or only the last layer. This observation further corroborates our assertion in "SMILES Alignment" section that supervised SMILES alignment methods might diminish the diversity of cross-attention maps across different layers, consequently impairing the model's capacity for representation. In contrast, unsupervised SMILES alignment methods do not exert such an adverse influence. This is why our unsupervised SMILES alignment mechanism achieves better results than supervised SMILES alignment.

Case study (multi-step retrosynthetic pathway planning)

To explore the suitability of our model for multi-step retrosynthetic pathway planning, we select three distinct molecules as targets for synthetic route design, and the synthesis routes are obtained through iterative calls to our UAlign model, which is trained with the USPTO-FULL dataset. The predicted pathways are summarized in Fig. 4.

The first case study involves Mitapivat, a compound approved for the treatment of hereditary hemolytic anemias in February 2022 [1]. Our model successfully predicted the five-step synthetic route reported in [3], with each step consistently ranked within the top-2 predictions. The first step entails an amide coupling reaction, which our model placed at rank 2, yielding the reactants 1-(cyclopropylmethyl)piperazine (compound 2) and 4-(quinoline-8-sulfonamido)benzoic acid (compound 3). Notably, at the initial step, our model also proposed an alternative synthesis method utilizing the Borch Reductive Amination, which was ranked at the first and is consistent with the synthetic route delineated by Saunders et al. Subsequently, for the synthesis of 4-(quinoline-8-sulfonamido)benzoic acid, the model precisely executed a functional group protection strategy during the second step and accurately anticipated the subsequent formation of the sulfonamide, effectively deconstructing the target molecule into readily available precursors.




Fig. 4 Multistep retrosynthesis predictions by our method. **a** Mitapivat **b** Pacritinib **c** Daprodustat. The reaction centers and leaving groups are highlighted in different colors. The pathway of molecules **a** and **b** come from literature, while the last one is verified by chemical experts

For the synthesis of 1-(cyclopropylmethyl)piperazine, the model strategically protected the amine functional group with a tert-butyloxycarbonyl moiety at the outset and, in the ultimate step, prognosticated the N-alkylation reaction with a top-ranking accuracy. This example illustrates our model's capability to uncover diverse reaction centers in molecular retrosynthetic design and to generate plausible reactant combinations based on these insights.

The second case under scrutiny is Pacritinib, an orally bioavailable and isoform selective JAK-2 inhibitor for the treatment of patients with myelofibrosis, which received FDA approval on February 28, 2022 [48]. As shown in Fig. 4(b), our model successfully delineates a eight-step synthesis, as described in the literature [4], tracing the synthetic pathway from commercially available 5-nitrosalicylaldehyde and 2,6-dichloropyrimidine to the final product. The initial step of the reverse synthesis is olefin

metathesis, ranking the first in order of likelihood, followed by another rank-2 aromatic substitution of 4-(3-((allyloxy)methyl)phenyl)-2-chloropyrimidine (compound **12**) and 3-((allyloxy)methyl)-4-(2-(pyrrolidin-1-yl)ethoxy)aniline (compound **13**). Subsequently, synthesis of 4-(3-((allyloxy)methyl)phenyl)-2-chloropyrimidine was correctly identified via continuous allyl substitution and Suzuki cross-coupling reaction as the top and the second choices. The reverse synthesis of 3-((allyloxy)methyl)-4-(2-(pyrrolidin-1-yl)ethoxy)aniline was reduction of the nitro group, followed by another allyl substitution. In the final step, the model's highest probability prediction was reduction of the aldehyde group, followed by a nucleophilic substitution. Despite the synthesis route involving a considerable number of steps and encompassing a variety of reaction types, our model successfully and accurately predicted each step within the top-2 choices. This accomplishment signifies the robustness and efficacy of our model in the context of retrosynthetic analysis.

The final case is Daprodustat, the first oral hypoxia-inducing factor prolyl hydroxylase inhibitor (HIF-PHI) for the treatment of renal anemia caused by chronic kidney disease (CKD) [11]. This novel compound received approval for market release from the FDA on the 1st of February, 2023 [48]. Our model predicted the three-step synthetic route. The first step reports the hydrolysis of ester at rank 3, which is aligned with the route provided by Duffy et al. Although next two steps provided by our method do not exist in the literature, there are all explainable. The synthesis of ethyl (1,3-dicyclohexyl-2,4,6-trioxohexahydropyrimidine-5-carbonyl)glycinate (compound **21**) was identified via dehydration condensation of 1,3-dicyclohexyl-2,4,6-trioxohexahydropyrimidine-5-carboxylic acid (compound **23**) and ethyl glycinate (compound **22**) as the top choice, which avoided using toxic ethyl isocyanatoacetate reported in literature. In the final step, the model's highest probability prediction was amidation of ester, resulting in cost-effective and readily accessible starting materials. This case demonstrates the robust extrapolative capacity of our model, highlighting its potential to generate synthetic routes that surpass those documented in the literature.

We also provide the results of multi-step retrosynthesis planning of two powerful baselines: SMILES-based method Retroformer [40] and graph-based method GraphRetro [33] in Supplementary Fig. 2 and Supplementary Fig. 3 respectively. The visualization reveals that while the Retroformer outperforms our method in the prediction of the synthetic route for Mitapivat, placing the literature pathway in a more advanced position, our model still accurately predicted each step of the literature-provided synthetic route within the top-2 choices. Conversely,

when faced with compounds like Pacritinib, which has multiple potential reaction centers, the Retroformer exhibits disadvantages. This is evident in Supplementary Fig. 2 (b) from steps 4 to 7, where the literature-documented synthetic route is ranked beyond the third position by Retroformer. Additionally, it is observable that Retroformer lacks robust predictive power for complex reactions, such as those requiring ring formation. Supplementary Fig. 2 illustrates that Retroformer failed to successfully predict Pacritinib's step-1 and was also unable to forecast both the literature pathways and the pathway validated by our chemical experts for Daprodust. Supplementary Fig. 3 further demonstrates that the performance of GraphRetro is marginally worse than that of Retroformer. Across the three presented cases, GraphRetro has not been able to successfully predict the synthetic routes. All the results above suggest that our model demonstrates a stronger capability in handling more complex molecules and reaction types compared to our baselines.

Discussion

We present UAlign, a novel graph-to-sequence pipeline that achieves state-of-the-art performance in the field of template-free methods. Our approach outperforms existing template-free and semi-template-based methods, while achieving comparable results to template-based methods. By utilizing a specially-designed graph neural network as the encoder, our model effectively leverages chemical and structural information from molecule graphs, resulting in powerful embedding for the decoder. Additionally, our proposed unsupervised SMILES alignment mechanism facilitates the reuse of shared substructures between reactants and products for generation, allowing the model to prioritize chemical knowledge even without complex data annotations. This significantly enhances the performance of the pipeline.

Despite achieving commendable performance, our work acknowledges areas for improvement. This work does not integrate much domain knowledge related to chemical reaction mechanisms in its design, which to some extent, compromises its interpretability. Similarly to most template-free methods, our work also faces challenges in generating diverse results. Additionally, we recognize a significant disparity between single-step retrosynthesis prediction and the complex reality of molecular synthesis route planning, underscoring the need for more realistic evaluation metrics to validate proposed models. To this end, we are charting a course for future exploration.

We remain committed to monitoring advancements in the understanding of chemical reaction mechanisms and intend to compile pertinent information from the field

of chemical reactions, including kinetic and thermodynamic data, to construct a more interpretable single-step retrosynthesis prediction model. We also plan to build a multiple-step route design system for molecular synthesis, integrating UAlign with search algorithms and predictive models for reaction conditions. Based on such a system, we expect to work with chemists to synthesize complex molecules in the wet lab, assessing the capability of the retrosynthesis prediction model.

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s13321-024-00877-2>.

Supplementary file 1.

Author contributions

K.Z. was responsible for the code implementation, algorithm design, and the majority of the manuscript writing. X.Z. conducted the Ablation Study. Y.Z. and B.Y. managed multi-step Retrosynthetic Pathway Planning. F.N. handled all visualizations. The remaining authors contributed to the polishing of the article. Y.X., Y.J. and X.Y. supervised the research.

Funding

This work was supported by the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102) and the Fundamental Research Funds for the Central Universities. The computations in this paper were run on the AI for Science Platform supported by the Artificial Intelligence Institute at Shanghai Jiao Tong University.

Data availability

The Data of USPTO-FULL and USPTO-50K can be found in <https://github.com/Hanjun-Dai/GLN>. The Data of USPTO-MIT can be found in <https://github.com/wengong-jin/nips17-rexgen/tree/master/USPTO>. We also provide the raw data and our processed version in <https://github.com/zengkaipeng/UAlign>.

Code availability

All the codes and checkpoints can be found in <https://github.com/zengkaipeng/UAlign>.

Materials availability

Not applicable.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no Competing interests.

Author details

¹MoE Key Laboratory of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai 200240, Shanghai, China. ²Frontiers Science Center for Transformative Molecules (FSCTM), Zhangjiang Institute for Advanced Study, Shanghai Jiao Tong University, Shanghai 200240, Shanghai, China. ³Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, Shanghai, China.

Received: 19 April 2024 Accepted: 30 June 2024

Published online: 15 July 2024

References

1. Al-Samkari H, van Beers EJ (2021) Mitapivat, a novel pyruvate kinase activator, for the treatment of hereditary hemolytic anemias. *Ther Adv Hematol* 12:20406207211066070. <https://doi.org/10.1177/20406207211066070>
2. Ba JL, Kiros JR, Hinton GE (2016) Layer normalization
3. Benedetto Tiz D, Bagnoli L, Rosati O et al (2022) Fda-approved small molecules in 2022: clinical uses and their synthesis. *Pharmaceutics* 14(11):2538
4. Chang H, Yajun G, Jiajuan P et al (2015) Synthesis of pacritinib hydrochloride. *Chin J Pharm*. <https://doi.org/10.16522/j.cnki.cjph.2015.12.001>
5. Chen S, Jung Y (2021) Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au* 1(10):1612–1620
6. Chen Z, Ayinde OR, Fuchs JR et al (2023) G 2 retro as a two-step graph generative models for retrosynthesis prediction. *Commun Chem*. <https://doi.org/10.1038/s42004-023-00897-3>
7. Coley CW, Rogers L, Green WH et al (2017) Computer-assisted retrosynthesis based on molecular similarity. *ACS Central Sci* 3(12):1237–1245
8. Corey EJ (1991) The logic of chemical synthesis: multistep synthesis of complex carbogenic molecules (nobel lecture). *Angew Chem Int Edn Eng* 30(5):455–465
9. Dai H, Li C, Coley C et al (2019) Retrosynthesis prediction with conditional graph logic network. *Advances in Neural Information Processing Systems* 32
10. Duffy K, Fitch D, Jin J et al (2007) Preparation of n-substituted pyrimidine-trione amino acid derivatives as prolyl hydroxylase inhibitors. WO2007150011A2
11. Hara K, Takahashi N, Wakamatsu A et al (2015) Pharmacokinetics, pharmacodynamics and safety of single, oral doses of gsk1278863, a novel hif-prolyl hydroxylase inhibitor, in healthy japanese and caucasian subjects. *Drug Metabolism and Pharmacokinetics* 30(6):410–418. <https://doi.org/10.1016/j.dmpk.2015.08.004>, <https://www.sciencedirect.com/science/article/pii/S1347436715000518>
12. Holtzman A, Buys J, Du L et al. The curious case of neural text degeneration. In: International Conference on Learning Representations, <https://openreview.net/forum?id=rygGQyrFvH>. 2020.
13. Hu W, Liu B, Gomes J et al. Strategies for pre-training graph neural networks. In: International Conference on Learning Representations. 2019.
14. Igashov I, Schneuing A, Segler M et al. Retrobridge: modeling retrosynthesis with markov bridges. In: The Twelfth International Conference on Learning Representations. 2023.
15. Ishida S, Miyazaki T, Sugaya Y et al (2021) Graph neural networks with multiple feature extraction paths for chemical property estimation. *Molecules* 26(11):3125
16. Jin W, Coley C, Barzilay R et al (2017) Predicting organic reaction outcomes with weisfeiler-lehman network. *Advances in neural information processing systems*. 2017; 30.
17. Kao YT, Wang SF, Wu MH et al (2022) A substructure-based screening approach to uncover n-nitrosamines in drug substances. *J Food Drug Anal* 30(1):150
18. Kim E, Lee D, Kwon Y et al (2021) Valid, plausible, and diverse retrosynthesis using tied two-way transformers with latent variables. *J Chem Inf Model* 61(1):123–133
19. Landrum G et al (2013) Rdkit: a software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum* 8:31
20. Lee J, Lee Y, Kim J et al (2019) Set transformer: A framework for attention-based permutation-invariant neural networks. In: Chaudhuri K, Salakhutdinov R (eds) Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol 97. PMLR, pp 3744–3753, <https://proceedings.mlr.press/v97/lee19d.html>
21. Li Y, Zhao H (2023) EM pre-training for multi-party dialogue response generation. In: Rogers A, Boyd-Graber J, Okazaki N (eds) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics, Toronto, Canada, pp 92–103, <https://doi.org/10.18653/v1/2023.acl-long.7>, <https://aclanthology.org/2023.acl-long.7>
22. Li Y, Huang X, Bi W et al (2023) Pre-training multi-party dialogue models with latent discourse inference. In: Rogers A, Boyd-Graber J, Okazaki N (eds) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for

- Computational Linguistics, Toronto, Canada, pp 9584–9599, <https://doi.org/10.18653/v1/2023.acl-long.533>, <https://aclanthology.org/2023.acl-long.533>
23. Lin K, Xu Y, Pei J et al (2020) Automatic retrosynthetic route planning using template-free models. *Chem Sci* 11(12):3355–3364
 24. Liu B, Ramsundar B, Kawthekar P et al (2017) Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS Central Sci* 3(10):1103–1113
 25. Radford A, Narasimhan K, Salimans T et al (2018) Improving language understanding by generative pre-training. OpenAI Technical Report https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
 26. Rong Y, Bian Y, Xu T et al (2020) Self-supervised graph transformer on large-scale molecular data. In: Larochelle H, Ranzato M, Hadsell R, et al (eds) *Advances in Neural Information Processing Systems*, vol 33. Curran Associates, Inc., pp 12559–12571, https://proceedings.neurips.cc/paper_files/paper/2020/file/94aef38441efa3380a3bed3faf1f9d5d-Paper.pdf
 27. Sacha M, Błaz M, Byrski P et al (2021) Molecule edit graph attention network: modeling chemical reactions as sequences of graph edits. *J Chem Inf Model* 61(7):3273–3284
 28. Saunders JO, Salituro FG, Yan S (2010) Preparation of aroylpiperazines and related compounds as pyruvate kinase m2 modulators useful in treatment of cancer. *US2010331307A1*, 30 Dec 2010
 29. Schwaller P, Laino T, Gaudin T et al (2019) Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS Central Sci* 5(9):1572–1583
 30. Segler MH, Waller MP (2017) Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry-A Eur J* 23(25):5966–5971
 31. Seo SW, Song YY, Yang JY et al. Gta: Graph truncated attention for retrosynthesis. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2021; 35(1):531–539. <https://doi.org/10.1609/aaai.v35i1.16131>, <https://ojs.aaai.org/index.php/AAAI/article/view/16131>
 32. Shi C, Xu M, Guo H et al. A graph to graphs framework for retrosynthesis prediction. In: *International conference on machine learning*, PMLR, pp 8818–8827. 2020.
 33. Somnath VR, Bunne C, Coley C et al (2021) Learning graph models for retrosynthesis prediction. *Adv Neural Inf Process Syst* 34:9405–9415
 34. Tetko IV, Karpov P, Van Deursen R et al (2020) State-of-the-art augmented nlp transformer models for direct and single-step retrosynthesis. *Nat Commun* 11(1):5575
 35. Tu Z, Coley CW (2022) Permutation invariant graph-to-sequence model for template-free retrosynthesis and reaction prediction. *J Chem Inf Modeling* 62(15):3503–3513
 36. Ucak UV, Ashyrmamatov I, Ko J et al (2022) Retrosynthetic reaction pathway prediction through neural machine translation of atomic environments. *Nat Commun* 13(1):1186
 37. Vaswani A, Shazeer N, Parmar N et al. Attention is all you need. *Advances in neural information processing systems*. 2017.
 38. Veličković P, Cucurull G, Casanova A et al. Graph attention networks. In: *International Conference on Learning Representations*. 2018.
 39. Vijayakumar AK, Cogswell M, Selvaraju RR et al (2016) Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*
 40. Wan Y, Hsieh CY, Liao B et al (2022) Retroformer: Pushing the limits of end-to-end retrosynthesis transformer. In: *International Conference on Machine Learning*, PMLR, pp 22475–22490. 2022.
 41. Wang X, Li Y, Qiu J et al (2021) Retroprime: a diverse, plausible and transformer-based method for single-step retrosynthesis predictions. *Chem Eng J*. <https://doi.org/10.1016/j.cej.2021.129845>
 42. Weininger D (1988) Smiles, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J Chem Inf Comput Sci*. 28(1):31–36
 43. Wu Q, Zhao W, Li Z et al (2022) Nodeformer: a scalable graph structure learning transformer for node classification. *Adv Neural Inf Process Syst* 35:27387–27401
 44. Xie S, Yan R, Guo J et al (2023) Retrosynthesis prediction with local template retrieval. *Proceedings of the AAAI Conference on Artificial Intelligence* 37(4):5330–5338. <https://doi.org/10.1609/aaai.v37i4.25664>, <https://ojs.aaai.org/index.php/AAAI/article/view/25664>
 45. Yan C, Ding Q, Zhao P et al (2020) Retroxpert: decompose retrosynthesis prediction like a chemist. *Adv Neural Inf Processing Syst* 33:11248–11258
 46. Yang N, Zeng K, Wu Q et al (2023) Molerec: combinatorial drug recommendation with substructure-aware molecular representation learning. *Proc ACM Web Conf* 2023:4075–4085
 47. Yao L, Guo W, Wang Z et al. Node-aligned graph-to-graph: elevating template-free deep learning approaches in single-step retrosynthesis. *JACS Au*. 2024.
 48. Zhang JY, Wang YT, Sun L et al (2023) Synthesis and clinical application of new drugs approved by fda in 2022. *Mol Biomed* 4(1):26
 49. Zheng S, Rao J, Zhang Z et al (2019) Predicting retrosynthetic reactions using self-corrected transformer neural networks. *J Chem Inf Modeling* 60(1):47–55
 50. Zhong Z, Song J, Feng Z et al (2022) Root-aligned smiles: a tight representation for chemical reaction prediction. *Chem Sci* 13(31):9023–9034

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.