LEARNING SINGLE-STEP RETROSYNTHESIS WITH SIMULATED REACTIONS

Shuan Chen[1] **and Yousung Jung**[2]
[1]KAIST, Daejeon, South Korea
[2]Seoul National University, Seoul, South Korea
yousung.jung@snu.ac.kr

ABSTRACT

Retrosynthesis analysis aims to design reaction pathways and intermediates for a target compound. Emerging works have been developed to automate this process by machine learning (ML) approaches, which greatly accelerate the process of synthesis pathway design. As data-driven approaches, ML models learn the synthetic pathways from existing reaction data. Although there are multiple synthesis pathways to synthesize one target product, one reaction product usually only has one corresponding reactant set in the training dataset. Therefore, existing models were trained by considering all the other reaction pathways as negative labels, which potentially includes a huge amount of false-negative data. In this work, we generate virtually validated simulated reactions by enumerating local reaction templates on the known reaction center and investigate the effect of training retrosynthesis models with these simulated reactions. We found that not only prediction accuracy but both prediction diversity and prediction confidence are also largely increased when training retrosynthesis models with simulated reactions. Specifically, the round-trip accuracy of top-5 prediction is increased from 86.2% to 90.8%, and the round-trip accuracy of predictions having output scores greater than 0.5 is increased from 93.5% to 96.4%. Moreover, the ratio of predictions showing output scores greater than 0.5 is increased from 6.0% to 26.5%. We also show that models trained with simulated reactions have a preference to predict more diverse synthesis pathways, including the reactions that are rarely seen in the training set.

1 INTRODUCTION

Retrosynthesis, referring to the process of decomposing a target molecule to its precursors, has been studied for decades since the first formalization by Corey (1991). Along with the advances of machine learning (ML)-based inverse molecular design, the number of ML-based retrosynthesis models, including single-step and multi-step retrosynthesis, is increasing rapidly to bring the virtually designed molecules into the real world (Coley et al. (2018); de Almeida et al. (2019)). While the total synthesis of most molecules requires multiple synthesis steps, the quality of the multi-step retrosynthesis model is highly dependent on the quality of the single-step retrosynthesis model since the multi-step model repeatedly applies single-step models to search the building blocks (Hassen et al. (2022)).

In general, single-step retrosynthesis models can be categorized into template-based (Coley et al. (2017); Segler & Waller (2017); Dai et al. (2019); Chen & Jung (2021); Seidl et al. (2022)) and template-free models (Liu et al. (2017); Tetko et al. (2020); Tu & Coley (2022); Zhong et al. (2022)) depending on using reaction templates or not. Template-based models classify which of the pre-defined reaction templates should be applied to a given product to obtain its precursors. For example, Dai et al. (2019) proposed Graph Logic Network (*GLN*) to maximize the likelihood of predicting the correct reaction template among a set of pre-defined reaction templates for a given product. Chen & Jung (2021) trained *LocalRetro* to learn the potential reaction center and the local reaction template to be applied to the given product. On the other hand, template-free models predict the precursors without pre-defining reaction templates. Most of the template-free models are based on the sequence

translation model, which "translates" the SMILES string of the given product to SMILES strings of its reactants (Liu et al. (2017); Tetko et al. (2020); Zhong et al. (2022)).

Although the existing models have shown promising results, most of the methods formulate the retrosynthesis problem as a one-to-one classification problem, which trains the model to select one of the pre-defined templates (template-based) or tokens (template-free). Given the fact that one product has more than one potential synthesis pathway, retrosynthesis prediction is naturally a one-to-many classification problem. As only one of the reactants is given for a single product in most of the reactions in the dataset, training a retrosynthesis model by one-to-many classification is a challenging task due to the lack of other existing synthesis pathways.

In this work, we generate virtually validated simulated reactions to fill the blank of potential retrosynthesis pathways in the reaction dataset using local reaction templates (LRTs, Chen & Jung (2021)) and the reaction outcome prediction model, and use these reactions to train the model to learn multiple feasible reaction templates at the same reaction center. By training LocalRetro (Chen & Jung (2021)) with simulated reactions, we show the potential direction of using simulated reactions to improve the diversity as well as the predictive performance of retrosynthetic models.

2 METHODS

2.1 GENERATE SIMULATED REACTIONS

For a given reaction, we enumerate all applicable LRTs found in the LRT library, which includes all the LRTs extracted from the reaction dataset, at the known reaction center to generate simulated reactions. To ensure these reactions are chemically plausible, we use *LocalTransform* (Chen & Jung (2022)) to filter out infeasible reactions if the predicted product is different from the original product. Considering the synthesizability of reactants, we only generate reactions using similar LRT, which yields at least one identical reactant in the original reaction.

For example, the C-C bond between two aromatic rings in Figure 1a is formed by Suzuki-coupling using boronic acid and bromide in the reaction dataset. By searching the LRT library derived from the reaction dataset, we found C-C bond between two aromatic rings can also be formed by similar LRT, either by changing the halide to chloride or iodine, or changing the boronic acid to a magnesium ion (1+) or trimethylstannyl group. With these similar LRTs, four simulated reactions that can potentially be used to synthesize the same product are generated.

2.2 LOCALRETRO VARIANTS

In this paper, we use the template-based retrosynthesis model *LocalRetro* (Chen & Jung (2021)) as a baseline model and train two variant models using simulated reactions. Because the score generated from the original LocalRetro is the probability of LRT applied on the predicted reaction center, the score does not physically reflect the feasibility of the predicted synthesis pathway. Here, we binarize the prediction score of each reactant set R of a given product O to a binary probability $p(R|O)$, which is the product of the probability of selecting a reaction center for a given product $p(o|O)$ and probability of local reaction template predicted at the reaction center $p(T|o)$.

$$p(R|O) = p(T|o) * p(o|O) \tag{1}$$

The training loss of new LocalRetro is the average of the loss of predicting reaction center $Loss_{RC}$ and the loss of predicting local reaction template $Loss_{LRT}$, where $Loss_{RC}$ is the binary cross entropy between labeled reaction center y_n and predicted probability of each atom or bond being reaction center p_n, and $Loss_{LRT}$ is the binary cross entropy between labeled LRT y_t and predicted probability of each LRT p_t applied at the known reaction center.

$$Loss_{RC} = \sum_{n=1}^{N} -(y_n \log(p_n) + (1 - y_n) \log(1 - p_n)) \tag{2}$$

$$Loss_{LRT} = \sum_{t=1}^{T} -(y_t \log(p_t) + (1 - y_t) \log(1 - p_t)) \tag{3}$$

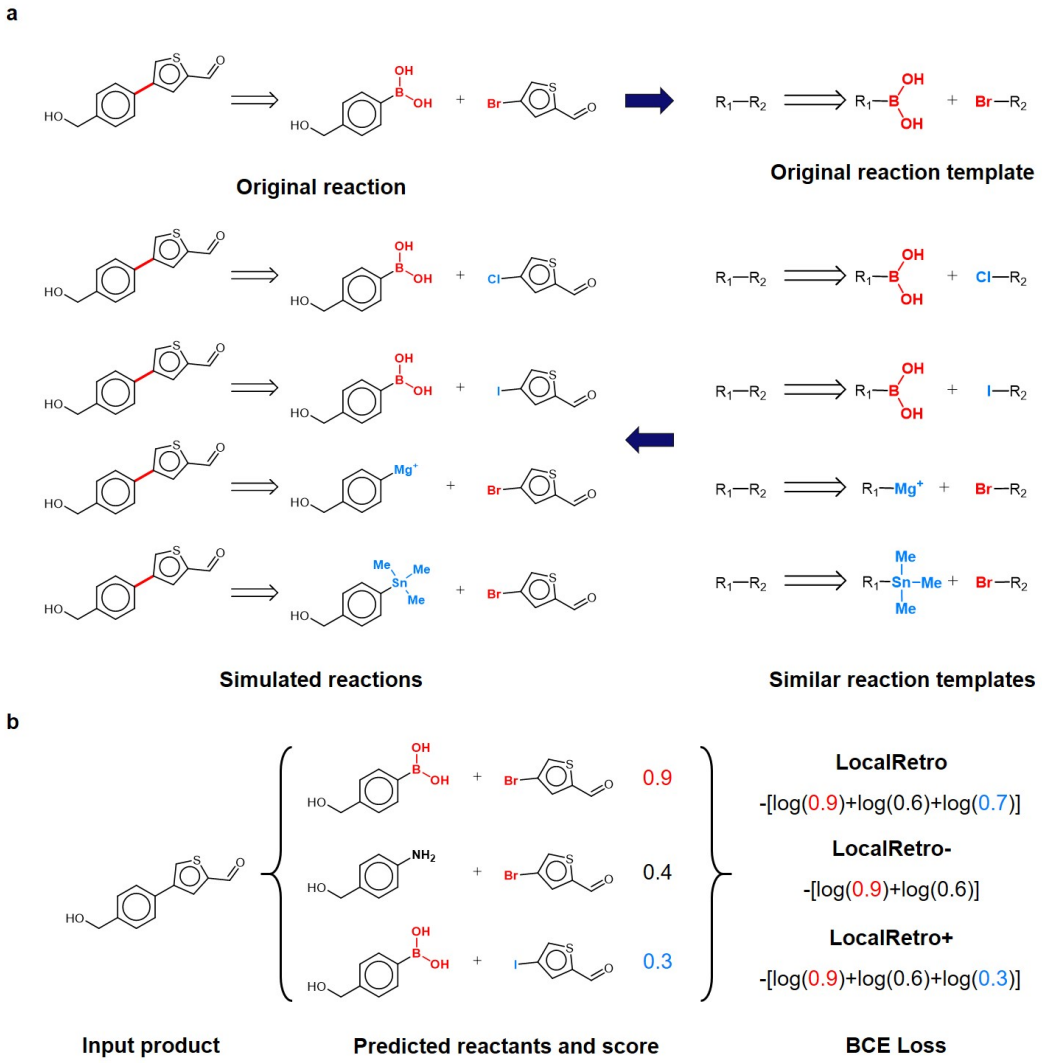

Figure 1: simulated reactions generation and the training loss of three different training settings. (a) The simulated reactions are generated by enumerating similar LRTs at the known reaction center, yielding at least one identical reactant with the original reaction. (b) Different training losses are used to optimize the model prediction in different LocalRetro variants. The loss from the original reaction is colored in red, the loss from the simulated reaction is colored in blue, and the loss from the infeasible reaction is colored in black.

Since we fixed the reaction center while generating simulated reactions, introducing simulated reactions will only change the LRT Loss ($Loss_{LRT}$). To investigate the effect of training LocalRetro with simulated reactions by either masking losses or inverting labels of simulated reactions to train the LocalRetro variants. When masking the losses from simulated reactions, the $Loss_{LRT}$ from simulated reactions are zeroed to reduce the penalties if the model predicts potentially feasible synthesis pathways. When inverting the labels of simulated reactions, the model is trained to maximize the prediction scores for both existing reactions and simulated reactions. In this work, we denote the former variant model as *LocalRetro-* and the latter variant model as *LocalRetro+* for reading convenience. An example of different training losses for LocalRetro and its variants is shown in Figure 1b.

3 EXPERIMENTS

3.1 DATASET

We evaluate our methods on the USPTO-50K dataset (Schneider et al. (2016)) with the same training/validation/test splits described in Coley et al. (2017). We extracted 678 LRTs from the 45K reactions in the training and validation set of the USPTO-50K dataset, including 92 atom templates and 586 bond templates. Total 123,098 simulated reactions are generated from the training and validation dataset.

3.2 EVALUATION METRIC

In this paper, we use round-trip accuracy (Schwaller et al. (2020)) to evaluate the feasibility of predicted reactants. Predictions are considered as correct if the products predicted by the reaction output prediction model match the original products. In this work, we use LocalTransform as the reaction outcome prediction for round-trip accuracy evaluation. We also calculate the conventional exact match accuracy of the proposed models in Supplementary B1.

3.3 MAIN RESULTS

We use the LocalRetro and its variants to predict 10 possible reactants for each product in the test set, and the round-trip accuracy at top-k predictions and the different threshold scores is shown in Table 1. For top-k predictions, the round-trip accuracy LocalRetro+ is improved by 0.6, 4.1, 4.6, 4.9% at k=1, 3, 5, 10 compared to the original LocalRetro. A smaller improvement is seen for LocalRetro- (2.2, 2.3, 2.6% at k=3, 5, 10). Next, we compare the round-trip accuracy of predictions having prediction scores over several threshold scores. Overall, the round-trip accuracy increases with the threshold for every model, and both LocalRetro- and LocalRetro+ show higher round-trip accuracy than the original LocalRetro at all the score threshold scores. Notably, the number of the predictions from LocalRetro- and LocalRetro+ over threshold scores 0.1, 0.5, and 0.9 are double and triple the number predicted from the original LocalRetro, respectively. These results support the training of retrosynthetic models with simulated reactions giving great improvements in both prediction accuracy and confidence.

Table 1: Round-trip accuracy on USPTO-50K dataset. The highest value is highlighted in bold font.

Model	**Round-trip accuracy (prediction ratio) %**			
Top-k	k=1	3	5	10
LocalRetro	91.4	88.0	86.2	82.0
LocalRetro-	91.4	90.2	88.5	84.6
LocalRetro+	**92.0**	**92.1**	**90.8**	**86.9**
Score	>0.9	>0.5	>0.1	>0.01
LocalRetro	94.2 (1.0)	93.5 (6.0)	89.4 (20.4)	86.6 (53.8)
LocalRetro-	96.7 (2.4)	95.8 (12.3)	**92.6** (35.6)	**89.1** (71.8)
LocalRetro+	**98.6 (3.8)**	**96.4 (26.5)**	92.1 **(63.8)**	**89.1 (89.8)**

To investigate the prediction diversity and confidence of proposed LocalRetro variants, we analyze the scores and template popularity in the training set of the top-10 predictions from LocalRetro, LocalRetro- and LocalRetro+ at Figure 2. Because templates showing low popularity in the training set only appears a few times during the training process, original LocalRetro rarely predicts LRTs having low popularity in the training set. Furthermore, due to a large number of negative labels in the training set, nearly 80% of the predictions are predicted with output scores lower than 0.1. When masking the false-negative loss from simulated reactions (LocalRetro-), predictions show overall higher scores. When inverting the negative labels of simulated reactions (LocalRetro+), the output scores of predictions are greatly boosted. More importantly, LocalRetro+ shows a higher preference to predict diverse LRTs regardless of the template popularity in the training set.

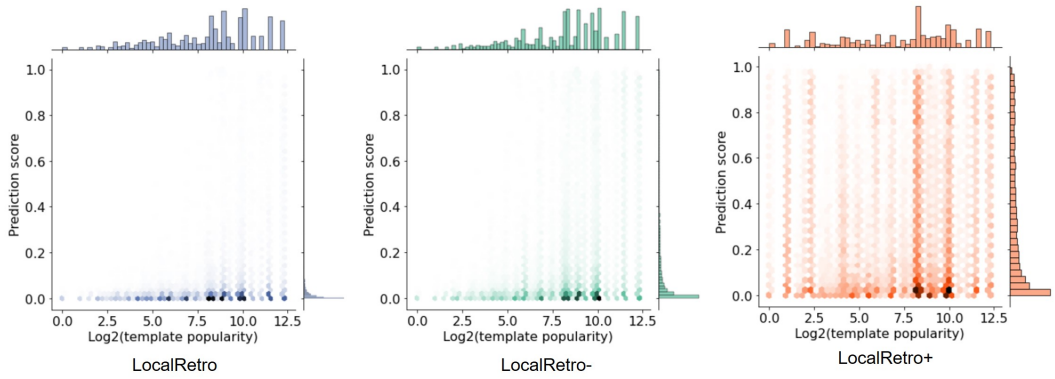

Figure 2: Scores and template popularity of predictions from LocalRetro variants.

3.4 Prediction examples

In this section, we qualitatively understand the difference between LocalRetro and its variant by showing an example of predicting the single-step retrosynthesis pathway of 2-methylbiphenyl in Figure 3. All of the LocalRetro models successfully predict the existing reactant in the dataset at top-1 prediction, and the remaining predictions are predicted to yield the same target product validated by LocalTransform. The first 5 predictions of LocalRetro and LocalRetro- are identical, while LocalRetro- gives a much higher prediction score for each prediction. Interestingly, while the top-5 predictions from LocalRetro and LocalRetro- suggest synthesizing the C-C bond between aromatic rings by Suzuki cross-coupling, LocalRetro+ suggests the synthesis pathways through cross-electrophile coupling, which is rarely seen in the USPTO-50K dataset.

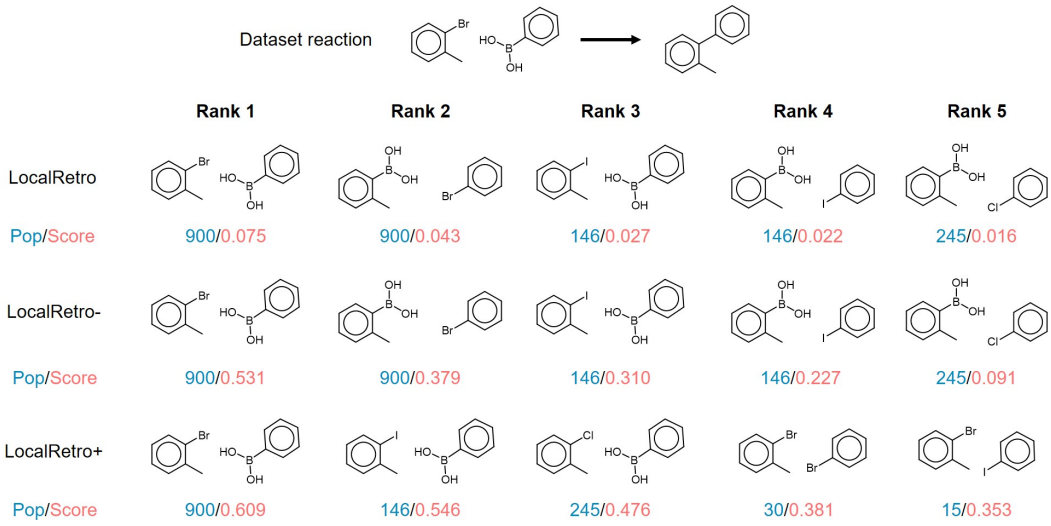

Figure 3: Example of predicting the single-step retrosynthetic pathway of 2-methylbiphenyl. The template popularity of the predicted LRT in the training set (Pop) and the prediction scores (Score) are displayed below the predicted reactant structures.

4 Conclusion

In this paper, we present two LocalRetro variants, LocalRetro- and LocalRetro+, to learn one-step retrosynthesis with simulated reactions by masking false-negative losses or inverting false-negative

labels. We found that both models give higher round-trip accuracy and prediction scores than the original LocalRetro. LocalRetro+, in particular, shows a higher preference to predict more diverse LRTs regardless of the popularity in the training set. Although we generate simulated reactions by applying all LRTs in the training set in this work, this method can be easily customized by selecting a few important LRTs based on the chemists' preferences. We suggest future retrosynthesis models may consider non-existent but potentially feasible synthesis pathways to train their models without extra actual experiments to prevent the model from learning false-negative reaction data.

AUTHOR CONTRIBUTIONS

S.C. and Y.J. conceived the project. S.C. designed the methods and performed the computational experiments and analyses. S.C. and Y.J. discussed the results and wrote the manuscript. Y.J. supervised the project.

ACKNOWLEDGMENTS

We acknowledge the support from the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (IITP-2021-0-02068, Artificial Intelligence Innovation Hub).

REFERENCES

Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021.

Shuan Chen and Yousung Jung. A generalized-template-based graph neural network for accurate organic reactivity prediction. *Nature Machine Intelligence*, 4(9):772–780, 2022.

Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Computer-assisted retrosynthesis based on molecular similarity. *ACS central science*, 3(12):1237–1245, 2017.

Connor W Coley, William H Green, and Klavs F Jensen. Machine learning in computer-aided synthesis planning. *Accounts of chemical research*, 51(5):1281–1289, 2018.

Elias James Corey. The logic of chemical synthesis: multistep synthesis of complex carbogenic molecules (nobel lecture). *Angewandte Chemie International Edition in English*, 30(5):455–465, 1991.

Hanjun Dai, Chengtao Li, Connor Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. *Advances in Neural Information Processing Systems*, 32, 2019.

A Filipa de Almeida, Rui Moreira, and Tiago Rodrigues. Synthetic organic chemistry driven by artificial intelligence. *Nature Reviews Chemistry*, 3(10):589–604, 2019.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. PMLR, 2017.

Alan Kai Hassen, Paula Torren-Peraire, Samuel Genheden, Jonas Verhoeven, Mike Preuss, and Igor Tetko. Mind the retrosynthesis gap: Bridging the divide between single-step and multi-step retrosynthesis prediction. *arXiv preprint arXiv:2212.11809*, 2022.

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Greg Landrum. Rdkit documentation. *Release*, 1(1-79):4, 2013.

Mufei Li, Jinjing Zhou, Jiajing Hu, Wenxuan Fan, Yangkang Zhang, Yaxin Gu, and George Karypis. Dgl-lifesci: An open-source toolkit for deep learning on graphs in life science. *ACS Omega*, 2021.

Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS central science*, 3(10):1103–1113, 2017.

Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Icml*, 2010.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Nadine Schneider, Nikolaus Stiefl, and Gregory A Landrum. What's what: The (nearly) definitive guide to reaction role assignment. *Journal of chemical information and modeling*, 56(12):2336–2346, 2016.

Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical science*, 11(12):3316–3325, 2020.

Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry–A European Journal*, 23(25):5966–5971, 2017.

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. *Journal of chemical information and modeling*, 62(9):2111–2120, 2022.

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. *Nature communications*, 11(1): 1–11, 2020.

Zhengkai Tu and Connor W Coley. Permutation invariant graph-to-sequence model for template-free retrosynthesis and reaction prediction. *Journal of chemical information and modeling*, 62 (15):3503–3513, 2022.

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.

Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Min Wu, Tingjun Hou, and Mingli Song. Root-aligned smiles: a tight representation for chemical reaction prediction. *Chemical Science*, 13(31):9023–9034, 2022.

A Model Details

A.1 LocalRetro

The LocalRetro model used in this paper is modified from the GitHub page of Chen & Jung (2021), which is a graph-based model using message passing neural network (MPNN, Gilmer et al. (2017) for local reactivity and attention mechanism (Vaswani et al. (2017)) for global reactivity.

First, all the ReLU activation functions (Nair & Hinton (2010)) are changed to GELU activation functions (Hendrycks & Gimpel (2016)), where $\Phi(x)$ is the cumulative distribution function of Gaussian distribution.

$$GELU(x) = x\Phi(x) \tag{4}$$

Second, the updated bond features are concatenated with their original bond features to avoid the vanishing of bond feature over MPNN, where v'_a and v'_b is the updated atom features after MPNN, e_{ab} is the original features of the chemical bond.

$$e'_{ab} = MLP(v'_a||v'_b||e_{ab}) \tag{5}$$

Third, a binary classifier of classifying whether an atom or bond can be the reaction center o of a given product molecule O, where v'' is the atom or bond features after a global reactivity attention block.

$$p(o|O) = MLP(v''_o) \tag{6}$$

B EXPERIMENTAL DETAILS

B.1 EXACT MATCH ACCURACY

The top-k exact match accuracy of LocalRetro and its variants tested on USPTO-50K is shown in Table 2. LocalRetro- shows an 0.1, 0.7, 0.3% accuracy improvement compared to the original model at top-5, 10, and 50 predictions. On the other hand, LocalRetro+ shows lower top-k accuracy across all k values. As LocalRetro+ has a higher preference to predict LRTs with low template popularity, LRTs with low template popularity in the training set are also not popular in the test set, leading to the drop in exact match accuracy.

Table 2: Top-k exact match accuracy on USPTO-50K dataset. The highest accuracy is highlighted in bold font.

Model	**Top-k accuracy %**				
	k=1	3	5	10	50
LocalRetro	**54.3**	**78.7**	**86.5**	**93.2**	**98.4**
LocalRetro-	50.9	73.8	83.2	92.0	**98.4**
LocalRetro+	36.4	57.9	69.6	82.9	97.8

B.2 IMPLEMENTATION

We use the open-source Python library RDKit (Landrum (2013)) for all the molecular pre- and post-processing. The machine learning models are highly dependent on dgllife (Li et al. (2021)) and Pytorch (Paszke et al. (2019)) deep learning libraries. All experiments were conducted under a machine equipped with an Intel Core i9-12900K @3.20 GHz, 128 GB of RAM, and NVIDIA GeForce RTX 3090 GPU.

B.3 TRAINING CONDITIONS

We train our models for 30 epochs with a batch size of 16 by using the Adam optimizer (Kingma & Ba (2014)) with a weight decay of 10^{-6} and set the initial learning rate to 10^{-4}. The learning rate is reduced by a factor of 0.1 after every 8 epochs and the training is terminated if the validation loss does not decrease for three successive epochs. Model gradients are clipped at a maximum norm of 20.