Fast prediction of distances between synthetic routes with deep learning

Samuel Genheden*, Ola Engkvist, Esben Bjerrum

Molecular AI, Discovery Sciences, R&D, AstraZeneca Gothenburg, SE-431 83 Mölndal, Sweden

* Corresponding author: samuel.genheden@astrazeneca.com

Abstract

We expand our recent work on clustering of synthesis routes and train a deep learning model to predict the distances between arbitrary routes. The model is based on a long short-term memory (LSTM) representation of a synthesis route and is trained as a twin network to reproduce the tree edit distance (TED) between two routes. The ML approach is approximately two orders of magnitude faster than the TED approach and enables clustering many more routes from a retrosynthesis route prediction. The clusters have a high degree of similarity to the clusters given by the TED-based approach and are accordingly intuitive and explainable. We provide the developed model as open-source (<https://github.com/MolecularAI/route-distances>).

Introduction

Machine learning and deep learning have propelled the research on computer-aided synthesis planning (CASP) in the last decade (see [1] for a recent review). CASP is becoming an integral part of drug development where models can be used to for instance predict reaction conditions or how to synthesize molecules. Retrosynthetic analysis is a technique used to find a synthetic route for a compound, in which the compound is broken down into smaller and smaller precursors until those precursors are purchasable or it is already known how to synthesize them [2, 3]. This can be formalized in a computer program, and today there exist many such algorithms [see e.g. 4, 5, 6, 7, 8, 9, 10] and software [1] to do this. Searching for synthesis routes requires an effective search algorithm as there are typically many ways to break down a compound into smaller building blocks. A search algorithm should not only return the synthesis route that is most likely to succeed, but additionally a set of diverse routes that a chemist can analyze further if the top-ranked route is not adequate. The pruning of the identified routes to create a diverse set of routes can be done as part of the search algorithm itself [11] or as a post-processing step [12].

We have recently proposed a clustering algorithm for predicted synthesis routes that is based on a tree edit distance (TED) calculation [13]. This is a graph theoretical metric [14, 15] of the distance between two synthesis routes that recursively utilizes chemical similarity matching. Our clustering approach can be used as a post-processing step to group together the results of a retrosynthetic analysis in order to aid the chemists in the analysis of the proposed routes. This algorithm could also be used to potentially compare predicted routes to reported experimental routes of the same compound, or to cluster different compounds based on the similarity of their synthesis routes. Although the TED calculations produce intuitive and explainable clusters, the method is unfortunately sometimes slow and does not scale to larger sets of routes [13].

Therefore, we decided to develop a machine learning (ML) model that is capable of predicting the distance between two arbitrary synthesis routes. This model is similar in design to the one

developed by Mo et al. to distinguish between predicted and experimental routes [12]. Although their model can be used to cluster synthesis routes, we observed some non-intuitive behavior of the model, which we believe is due to the training objective of the model. The model of Mo et al., was never trained for the distance calculation task, thus it is not guaranteed to perform well on this task. The latent space representation of the synthesis route that is calculated by the model is simply a by-product of the training task. On the other hand, the TED-based approach leads to qualitative good clusters [13] and can therefore serve as the ground truth for the training of an ML model. We used a model inspired by the model of Mo et al., however we trained it to reproduce the TED calculations. The ML model serves as a proxy model for the expensive TED calculations. Thus, a fast model can be developed that also leads to intuitive and explainable clusters.

Methods

Compound selection and preparation. We used a set of 5,000 compounds from ChEMBL [16], described and used previously for benchmarking the TED-based clustering [13]. This set will be referred to as ChEMBL-5k. We also created a new set by randomly sampling 10,000 single molecule compounds from ChEMBL that has a molecular weight between 100 and 800 D, a QED [17] score above 0.2 and is not part of the ChEMBL-5k set. The tautomeric form of the compounds was determined by RDKit [18]. This set will be referred to as ChEMBL-10k. Finally we also randomly sampled 10,000 compounds from the GDB ChEMBL and 10,000 compounds from the GDB MedChem datasets [19, 20]. The structures provided by these sets were used without further processing. These sets will be referred to as GDB-ChEMBL and GDB-MedChem, respectively.

Route predictions and TED calculations. The compounds in the four different sets were subjected to retrosynthesis predictions with the AiZynthFinder software [21]. The expansion and filter policies used by the software were derived from the USPTO dataset [22, 23], unless otherwise stated. Enamine building blocks and those available internally at AstraZeneca were used as stop criteria. The retrosynthesis search was stopped after 100 iterations, and between 5 and 25 routes were extracted, depending on the search score. For each target compound, the pairwise distance matrix of the predicted routes was calculated using a tree edit distance (TED) algorithm, as detailed previously [13].

Machine learning model. We designed a machine learning model to represent a synthesis route that is based on a child-sum tree LSTM (long short-term memory) neural network model [24, 25, 26]. On each molecular node in the route we position an LSTM cell, which takes input from children nodes (i.e. precursor molecules) as well as a feature representation of the molecule. In the forward pass of the model, the latent representation of the LSTM cells are updated iteratively until the top-node (the target molecule) is reached. The feature representation of a molecule is calculated by a simple feed-forward network that takes as input a 2048 bit fingerprint (ECFP4 [27], computed by the Morgan algorithm in RDKit [18]) of the molecule. A representation of the complete network architecture is shown in Figure 1. This model is similar to the model proposed by Mo et al. [12], although we do not consider reaction fingerprints and we use a different activation function in the feed-forward network.

ML model training. We trained a twin network [28] based on the LSTM-model described above on the TED routes. Given a pair of routes, the first route is fed through the LSTM-model giving the latent representation of the top-node, followed by the feeding of the second route, also giving a latent representation of the top-node (see Figure 1). The Euclidean distance

between these two latent vectors should reproduce the TED results. We used a mean squared error-loss function together with the mean absolute error metric to monitor the training. The Adam optimizer [29] was used with a learning rate of 0.001 and a weight decay factor of 0.001 [30]. The learning rate was decreased after a relative plateau of 10^{-4} of the loss was observed after 10 epochs. The dropout probability for the feed-forward network was 0.4 and the size of the LSTM latent space was 1024. We trained in batches of 128 samples for 50 epochs. Limited hyperparameter optimization was carried out with the Optuna package [31], but because the optimization is expensive and because good performance was obtained with the above parameters, we did not attempt to fine-tune them further. The model was implemented in PyTorch [32] using the PyTorch Lightning framework [33].

Figure 1 – Illustration of the machine learning model. A) A synthetic route consisting of two reactions and five molecules, B) The conceptual representation of the route by a LSTM network. A LSTM cell is placed on top of each molecule, which takes input from the children nodes as well as a processed fingerprint of the molecule. C) The architecture of the compression feed-forward takes a fingerprint (with length FP size) as input and produces a shortened feature representation of the molecule (with length LSTM size). D) The architecture of the twin network that is used to predicted distances between routes. For each route, the molecular fingerprints are fed through the compression network and the output representation is fed together with an adjacency matrix to a tree LSTM network. The adjacency matrix determines the node order in the iterative procedure that updates the LSTM cells. The LSTM encodings of the top nodes from the two routes are then used to compute an Euclidean distance.

The dataset consists of pairs of routes (T_i, T_j) and the calculated TED. Either both permutations of the pairs were included ($i < j, i = j, i > j$) or we only included one permutation of the pair ($i \leq j$). The dataset was split into approximately 80% for training, 10% for validation and 10% for testing, respectively. Care was taken so that a route is not in more than one set: we iteratively selected all pairs of routes for a random compound until we had a training set consisting of at least 80% of the total set of pairs. This procedure was repeated until we had selected at least 10% of the total set of pairs for validation and finally, the approximately 10% of pairs left were taken as the test set. As all pairs of routes originate from the same target compound, this ensures that a route is not in more than one set. Models were trained on routes generated for the

ChEMBL-5k or ChEMBL-10k compound sets. The number of routes and number of pairs in the different sets is shown in Table 1.

Clustering. Clustering was carried out based on the distance matrix from either TED calculations or the predictions of the ML model. We used hierarchical clustering with single linkage as implemented in SciKit-Learn [34], and the optimal number of clusters was determined by the Silhouette method [35]. The maximum number of clusters was set to five, because the number of analyzed routes were small in order to enable comparison between the ML and TED approaches.

Table 1 – The total number of routes and pairs in the different sets used to train the LSTM model.

Compound set	Pair permutations	# routes	# pairs
ChEMBL-10k	Single	104120	751666
ChEMBL-10k	All	104120	1399212
ChEMBL-5k	Single	51694	371175
ChEMBL-5k	All	51694	690656

Results and discussion

Figure 2 - Average validation loss, MAE, and R over three independent modelling runs for ChEMBL-5k and ChEMBL-10k datasets (single pair permutation). The grey area indicate an error bar of one standard deviation.

The ML model learns quickly. We trained an ML model to reproduce tree edit distances (TED) using mainly two different datasets: single permutation of pairs for ChEMBL-5k compounds, and single permutation of pairs for ChEMBL-10k compounds. The total number of routes and pairs is listed in Table 1. The average of the loss, the mean absolute error, and R metrics over three independent training runs are shown in Figure 2. All metrics converge after about 30 epochs, and there seems to be only a small variation between the different training runs. Therefore, we proceed with using the models produced from the first training runs when evaluating them further below. In Table S1, we compare the MAE and R when using a single permutation of the pairs to using all permutations of the pairs. For both MAE and R , the models trained on all pairs perform slightly worse than the models trained on only single permutations, although it is unclear if the differences are significant due to the low number of independent training runs. Therefore, we continued evaluation on the models trained on single pair permutations, because we can speed up the training time by only including about half the number of pairs. As the twin network share weights for the two LSTM models and because the Euclidean distance calculation is commutative, it should not matter which route is passed first through the network, which furthers motivated us to continue with the single-permutation training sets. Finally, there appears to be no significant difference in performance when comparing the ChEMBL-5k and ChEMBL-10k models.

Table 2 – Timings in seconds for TED and ML-based clustering for four datasets

	Mean # pairs	TED		ML	
		mean time	worst time	mean time	worst time
ChEMBL-5k	63.90	6.45	213.95	0.05	0.70
ChEMBL-10k	64.75	5.54	211.89	0.05	1.41
GDB-MedChem	50.76	6.71	156.83	0.06	2.24
GDB-ChEMBL	52.97	6.82	157.30	0.06	2.16
ChEMBL-5k (100 routes)	4770.86	377.03	4215.92	0.72	5.62
ChEMBL-5k (all solved routes)	1217.81	72.62	3806.78	0.20	4.65

ML-based clustering is significantly faster than TED-based clustering. The ML model trained on the ChEMBL-10k dataset (single pair permutation) was used to predict distances between routes for the ChEMBL-5k, GDB-MedChem, and GDB-ChEMBL compound sets. For the ChEMBL-10k compound sets, the model trained on the ChEMBL-5k dataset was used. The foremost reason to develop a proxy model such as the ML model is that the TED calculations were observed to be too slow in the worst case and did not scale well. Therefore, it is of interest to compare the timings of the two approaches. The mean and worst clustering time (distance + clustering calculations) is shown in Table 2. If we only look at the cases where we extract a small number of routes (first four rows), the mean time for TED calculation is between approximately 6 and 7 seconds, compared to 0.1 for the ML approach. The worst time is between 157 and 214 seconds for the TED approach and only between 0.7 and 2 seconds for the ML approach. This shows that both on average and in worst-case scenario the ML approach is more than an order of magnitude faster than the TED calculations, and potentially two orders faster. If we instead extract 100 routes for each compound, we see that the speed-up is even more impressive. The mean and worst time is only 0.7 and 6 s, respectively, which is negligible compared to the search time. In addition to extracting 100 routes, we also extract all solved

routes for compounds where a solution could be found, and between five and 25 routes for compounds where a solution could not be found. These timings are shown in Table 2 as well, and we can conclude that the average and worst clustering time is acceptable for this scenario as well. Therefore, we can conclude that we have successfully created a faster proxy model.

Table 3 – Metrics show the similarity of distances and clusters when comparing the TED and ML approaches for various datasets. The ML model was trained on ChEMBL-10K

Compound set	Expansion policy	Distance		Cluster	
		R	MAE	Mean similarity	Median similarity
ChEMBL-5k	USPTO	0.95	0.92	0.88	0.97
ChEMBL-10k	USPTO	0.95	1.03	0.87	0.95
GDB-MedChem	USPTO	0.92	1.66	0.87	0.96
GDB-ChEMBL	USPTO	0.92	1.61	0.87	0.95
GDB-MedChem	Reaxys	0.92	1.55	0.88	1.00
GDB-ChEMBL	Reaxys	0.92	1.53	0.88	1.00

ML-based clustering recovers the TED clusters to a large extent. In Table 3 we show the mean absolute error (MAE) and correlation coefficient, R when comparing the TED with the distances predicted by the ML model. For the ChEMBL-5k set, R is 0.95 and the MAE is 1 distance unit, which shows that there is a strong correlation and only a small deviation. This can also be seen in Figure 3; there are a few distances that are considerably different, especially distances where the ML predictions is much lower than the TED. We see similar R and MAE if we compare the distances for the routes of the ChEMBL-10k set, showing that a model trained a smaller dataset (Chembl-5k) is as good as training on a larger dataset (Chembl-10k). Having such a good prediction of distances, it is not surprising that the clustering performance is good as well. The mean similarity is between 0.87 and 0.88, with the median being between 0.95 and 0.97. This shows that most of the routes that are clustered together with TED are also clustered together with the ML approach. We can thus conclude that the ML approach at least on average produces satisfactory clusters.

Figure 3 – Performance of the ML model trained on the ChEMBL-10k dataset and applied on routes for ChEMBL-5k compounds. a) Correlation between TED and ML prediction of distances, b) The distribution of cluster similarities for all compounds.

ML model for route distances and clustering is transferable. We included compounds from GDB-MedChem and GDB-ChEMBL in the study to investigate the transferability power of the model. The GDB database contains enumerated compounds that in general are harder to break

down with the AiZynthFinder tool (see Table S2). This shows that the compounds in GDB contain chemistry that is not well-represented in the USPTO dataset, on which the expansion policy was trained. And even if one of the GDB sets were created to be similar to ChEMBL compounds [20], the compounds in this set are clearly much harder to break down than the compounds ChEMBL-10k and ChEMBL-5k sets, indicating that the compounds are different at least from a synthesizability perspective. However, as we see in Table 2 the performance of the ML approach on the GDB sets is very satisfactory. The correlation coefficient is slightly weaker than for the ChEMBL sets, and the MAE points at a larger discrepancy. Still, the mean and median cluster similarities for the GDB sets are on a par with the similarities seen for the ChEMBL sets, showing that the clustering algorithm is robust to distance discrepancies. This can be understood by the fact that we use the same approach to generate routes for ChEMBL and GDB compounds, which would lead to the same type of reaction chemistry being encoded in the routes. Thus the relation between molecules in the route is not inherently different when comparing ChEMBL and GDB routes, and therefore the ML model can successfully transfer the knowledge from routes for ChEMBL compounds to routes for GDB compounds. We also predicted routes for the GDB sets using an expansion policy trained on the Reaxys database [22, 36], and we can see in Table S2 that this expansion policy is more successful than the USPTO-based expansion policy in finding routes for the GDB compounds. We still reproduce the TED calculations when we predict route distances and clusters using the ML model trained in ChEMBL-10k: as seen in Table 2, the R and MAE when comparing the distances are 0.92 and approximately 1.5, respectively, and the mean cluster similarity is 0.88 for both GDB sets. Thus, we can be relatively certain that the trained ML model is predictive even if one changes the expansion policy or have compounds that are not ChEMBL-like.

Figure 4 – Dendrograms formed from distances matrixes from A) TED calculations, b) ML predictions.

Large clustering differences between TED-based and ML-based clustering can be rationalized. We investigate the difference clustering further by inspecting the distances and clusters for the compound where the cluster similarity was the lowest (0.28) when analyzing

the ChEMBL-5k set. For this compound, the optimal number of clusters according to the Silhouette method for the TED matrix is 5, whereas for the latent space distance matrix it is 2. This can be understood from the dendrograms of the two distance matrices shown in Figure 4. For the ML predictions, route number 4 is so distant from the other routes that it is placed into a singleton cluster, and all the other routes in a second cluster. For the TED approach, the distances are more evenly distributed and it is therefore more preferential to form more clusters. Still, the routes clustered first in the hierarchical clustering would be very similar for the two approaches. For instance, route 0 and 1 would be clustered early and then joined with route 3 in both TED and ML-based clustering. Similarly, the cluster consisting of routes 5 and 9 would be joined with the cluster formed from routes 6 and 10. That route 4 is different to the other routes can be rationalized by realizing that this route is the only one that ends with a dehydration reaction (see Figure 5). So the ML model single out this feature much more than the TED approach. TED clustering places route 4 closer to routes 0 to 3, which are single-step routes similar to the second step of route 4. This analysis highlights that the TED and ML approaches sometimes focus on different chemistry in the analyzed routes. It should be pointed out the optimal number of clusters has a subjective element to it, so that the two methods gives different number of clusters is not inherently wrong. For 58% of the compounds in the ChEMBL-5k set, the optimal number of clusters according to the Silhouette method is equal when comparing the TED and ML approaches, for 24% of the compounds the ML approach leads to more clusters and in 19% the TED approach leads to more clusters.

Figure 5 – Route number 4 from the dendrograms in Figure 4. It is the only route ending with a dehydration reaction.

Conclusions

A novel method to rapidly compute the similarity between synthetic routes has been developed and can be used to, e.g., cluster results from a retrosynthesis analysis, cluster compounds based on route similarity, compare predicted and experimentally reported routes, or in the comparison of synthesis planning tools. We showed that the novel method can reproduce the intuitive but much slower synthetic route clustering based on tree edit distances (TED). Because the cluster similarity between the ML and TED approaches is high, the ML model will provide chemists with useful clusters. Furthermore, we showed that the novel method is fast: on average the predictions take less than one second and in the worst-case a few seconds, which is negligible compared to the route prediction time. Finally, we also showed that the trained ML model is transferable: is robust to changes in targeted chemical space and the expansion policy used in the retrosynthesis analysis. We thus envisage that the novel method presented herein will be useful in synthesis planning. The developed method is provided as open-source (<https://github.com/MolecularAI/route-distances>) and is available in the AiZynthFinder software [21].

Acknowledgement

Amol Thakkar is acknowledged for proof-reading the text.

References

- [1] Johansson S, Thakkar A, Kogej T, Bjerrum E, Genheden S, Bastys T, Kannas C, Schliep A, Chen H and Engkvist O 2020 AI-assisted synthesis prediction *Drug Discov. Today Technol.* **32-33**:65-72
- [2] Corey E J and Todd Wipke W 1969 Computer-assisted design of complex organic syntheses *Science*. **166** 178–92
- [2] Coley C W, Green W H and Jensen K F 2018 Machine Learning in Computer-Aided Synthesis Planning *Acc. Chem. Res.* **51** 1281–9
- [4] Heifets A and Jurisica I 2012 *Construction of New Medicines via Game Proof Search*. 1564-1570
- [5] Segler M H S, Preuss M and Waller M P 2018 Planning chemical syntheses with deep neural networks and symbolic AI *Nature* **555** 604–10
- [6] Klucznik T, Mikulak-Klucznik B, McCormack M P, Lima H, Szymkuć S, Bhowmick M, Molga K, Zhou Y, Rickershauser L, Gajewska E P, Toutchkine A, Dittwald P, Startek M P, Kirkovits G J, Roszak R, Adamski A, Sieredzińska B, Mrksich M, Trice S L J and Grzybowski B A 2018 Efficient Syntheses of Diverse, Medicinally Relevant Targets Planned by Computer and Executed in the Laboratory *Chem* **4** 522–32
- [7] Coley C W, Thomas D A, Lummiss J A M, Jaworski J N, Breen C P, Schultz V, Hart T, Fishman J S, Rogers L, Gao H, Hicklin R W, Plehiers P P, Byington J, Piotti J S, Green W H, John Hart A, Jamison T F and Jensen K F 2019 A robotic platform for flow synthesis of organic compounds informed by AI planning *Science* **365** eaax1566
- [8] Lin K, Xu Y, Pei J and Lai L 2020 Automatic retrosynthetic route planning using template-free models *Chem. Sci.* **11** 3355–64
- [9] Schwaller P, Petraglia R, Zullo V, Nair V H, Haeuselmann R A, Pisoni R, Bekas C, iuliano anna and Laino T 2020 Predicting Retrosynthetic Pathways using Transformer-Based Models and a Hyper-Graph Exploration Strategy *Chem. Sci.*
- [10] Chen B, Li C, Dai H and Song L 2020 Retro*: Learning Retrosynthetic Planning with Neural Guided A* Search. *arXiv*.arXiv:2006.15820
- [11] Shibukawa R, Ishida S, Yoshizoe K, Wasa K, Takasu K, Okuno Y, Terayama K and Tsuda K 2020 CompRet: A comprehensive recommendation framework for chemical synthesis planning with algorithmic enumeration *J. Cheminform.* **12** 52
- [12] Mo Y, Guan Y, Verma P, Guo J, Fortunato M E, Lu Z, Coley C W and Jensen K F 2021 Evaluating and clustering retrosynthesis pathways with learned strategy *Chem. Sci.* **12** 1469-1478
- [13] Genheden S, Engkvist O, Bjerrum E. 2020 Clustering of Synthetic Routes Using Tree Edit Distance. ChemRxiv. Preprint. <https://doi.org/10.26434/chemrxiv.13372475.v1>
- [14] Pawlik M and Augsten N 2015 Efficient computation of the tree edit distance *ACM Trans. Database Syst.* **40** 1–40
- [15] Pawlik M and Augsten N 2016 Tree edit distance: Robust and memory-efficient *Inf. Syst.* **56** 157–73
- [16] Gaulton A, Bellis L J, Bento A P, Chambers J, Davies M, Hersey A, Light Y, McGlinchey S, Michalovich D, Al-Lazikani B and Overington J P 2012 ChEMBL: A large-scale bioactivity database for drug discovery *Nucleic Acids Res.* **40** D1100
- [17] Bickerton G R, Paolini G V., Besnard J, Muresan S and Hopkins A L 2012 Quantifying the chemical beauty of drugs *Nat. Chem.* **4** 90–8
- [18] RDKit: Open-source cheminformatics, <http://www.rdkit.org>.
- [19] Awale M, Sirockin F, Stiefl N and Reymond J L 2019 Medicinal Chemistry Aware Database GDBMedChem *Mol. Inform.* **38** 1900031
- [20] Bühlmann S and Reymond J L 2020 ChEMBL-Likeness Score and Database GDBChEMBL *Front. Chem.* **8** 46
- [21] Genheden S, Thakkar A, Chadimová V, Reymond J L, Engkvist O and Bjerrum E 2020 AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning *J. Cheminform.* **12** 70
- [22] Thakkar A, Kogej T, Reymond J L, Engkvist O and Bjerrum E J 2020 Datasets and their influence on the development of computer assisted synthesis planning tools in the pharmaceutical domain *Chem. Sci.* **11** 154–68
- [23] Genheden S, Engkvist O and Bjerrum E J 2020 A Quick Policy to Filter Reactions Based on Feasibility in AI-Guided Retrosynthetic Planning. *ChemRxiv*. Preprint.
- [24] Hochreiter S and Schmidhuber J 1997 Long Short-Term Memory *Neural Comput.* **9** 1735–80

- [25] Tai K S, Socher R and Manning C D 2015 Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks *ACL-IJCNLP 2015 - 53rd Annu. Meet. Assoc. Comput. Linguist. 7th Int. Jt. Conf. Nat. Lang. Process. Asian Fed. Nat. Lang. Process. Proc. Conf.* **1** 1556–66
- [26] <https://github.com/unbounce/pytorch-tree-lstm>
- [27] Rogers D and Hahn M. 2010 Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **50** 742–754
- [28] Chicco D 2021 Siamese Neural Networks: An Overview *Methods in Molecular Biology* vol 2190 (Humana Press Inc.) pp 73–94
- [29] Kingma D P and Ba J L 2015 Adam: A method for stochastic optimization *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings* (International Conference on Learning Representations, ICLR)
- [30] Loshchilov I and Hutter F 2017 Decoupled Weight Decay Regularization *arXiv* arXiv:1711.05101
- [31] Akiba T, Sano S, Yanase T, Ohta T and Koyama M 2019 Optuna: A Next-generation Hyperparameter Optimization Framework *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.* 2623–31
- [32] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Köpf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J and Chintala S 2019 PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXiv:1912.01703
- [33] Falcon W, Borovec J, Wälchli A, Eggert N, Schock J, Jordan J, Skafte N, Ir1dXD, Berezhnyuk V, Harris E, Murrell T, Yu P, Præsius S, Addair T, Zhong J, Lipin D, Uchida S, Bapat S, Schröter H, Dayma B, Karnachev A, Kulkarni A, Komatsu S, Martin B, Schiratti J-B, Mary H, Byrne D, Eyzaguirre C, cinjon and Bakhtin A 2020 PyTorchLightning/pytorch-lightning: 0.7.6 release. <https://github.com/PyTorchLightning/pytorch-lightning>
- [34] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E 2011 *Scikit-learn: Machine Learning in Python.* 12 2825–2830
- [35] Rousseeuw P J 1987 Silhouettes: A graphical aid to the interpretation and validation of cluster analysis *J. Comput. Appl. Math.* **20** 53–65
- [36] Reaxys©, Copyright © 2019 Elsevier Limited except certain content provided by third parties, Reaxys is a trademark of Elsevier

Supporting Information

Table S1 – Comparison of validation MAE and *R* when using either all pair permutations or only single permutations.

Compound set	Pair permutations	MAE	<i>R</i>
ChEMBL-5k	All	0.92±0.04	0.91±0.001
ChEMBL-5k	Single	0.84±0.02	0.93±0.004
ChEMBL-10k	All	0.87±0.02	0.92±0.002
ChEMBL-10k	Single	0.83±0.03	0.94±0.003

Table S2 – Statistics of the tree search. A compound is solved if at least one route is found where all the starting material is in stock.

Compound set	Expansion policy	# solved compounds	% of compounds sold	Mean search time
ChEMBL-5k	USPTO	3762	75.24	77.81
ChEMBL-10k	USPTO	7288	72.88	65.28
GDB-MedChem	USPTO	1882	18.82	91.75
GDB-ChEMBL	USPTO	2255	22.55	102.00
GDB-MedChem	Reaxys	3731	37.31	103.54
GDB-ChEMBL	Reaxys	3994	39.94	110.42