

TANGO*: CONSTRAINED SYNTHESIS PLANNING USING CHEMICALLY INFORMED VALUE FUNCTIONS

Daniel Armstrong, Zlatko Jončev, Jeff Guo & Philippe Schwaller*

Institute of Chemical Sciences and Engineering
Ecole Polytechnique Fédérale de Lausanne (EPFL)
1015 Lausanne, Switzerland
{daniel.armstrong, philippe.schwaller}@epfl.ch

ABSTRACT

Computer-aided synthesis planning (CASP) has made significant strides in generating retrosynthetic pathways for simple molecules in a non-constrained fashion. Recent work introduces a specialised *bidirectional* search algorithm with forward and retro expansion to address the starting material-constrained synthesis problem, allowing CASP systems to provide synthesis pathways from specified starting materials, such as waste products or renewable feed-stocks. In this work, we introduce a simple guided search that allows us to solve the starting material-constrained synthesis planning problem using an existing unidirectional search algorithm, Retro*. We show that by optimising a single hyperparameter, Tango* outperforms existing methods in terms of efficiency and solve rate. We find that the Tango* cost function catalyses strong improvements for the bidirectional DESP methods. Our method also achieves lower wall clock times while proposing synthetic routes of similar length, a common metric for route quality.

1 INTRODUCTION

Synthesis planning, where chemists design routes of chemical reactions to synthesise a complex molecule from simple or purchasable building blocks, is a key task in synthetic chemistry. The process used for this, *retrosynthetic analysis*, involves recursively performing *reversed* reactions, where a bond is broken to simplify a molecule into two or more component precursors (Corey, 1967; Corey & Chelg, 1995). Originally proposed by Corey in 1969, Computer-Assisted Synthesis Planning (CASP) aims to automate this process (Corey & Wipke, 1969). Since the seminal patent mining work of Lowe, which provided a large dataset of machine-readable chemical reactions, the CASP field has expanded significantly, with a plethora of approaches developed (Schneider et al., 2016; Liu et al., 2017; Segler & Waller, 2017; Coley et al., 2017; Chen et al., 2020; Tu & Coley, 2021; Sacha et al., 2021). CASP systems typically have two primary components: a single-step retrosynthesis model, which decomposes a molecule into simpler precursors, and a search algorithm that explores the search graph constructed from outputs of the single-step model (Segler & Waller, 2017; Segler et al., 2018; Schwaller et al., 2020; Chen et al., 2020; Kishimoto et al.; Genheden et al., 2020). The iterative application of single-step models and exploration of the generated search space typically continues until a molecule is "solved," which is specified as having all leaf nodes belonging to a predefined set of purchasable building blocks. This approach of finding a path to any available precursor differs substantially from the approach expert chemists may take, where chemists can plan a synthesis with numerous constraints in mind, such as avoiding certain reactions and solvents, or starting from a specific precursor, known as a "structure-goal" (Corey & Chelg, 1995). By starting from a building block containing a key structural motif, the overall molecular complexity gain in a synthesis route can be reduced, a technique called "semi-synthesis" (Ojima et al., 1992; Brill et al., 2017). There is also considerable interest in the recycling of waste compounds into useful products, a technique called "waste valorisation" (Wołos et al., 2022; Lopez et al., 2024; Żadło-Dobrowolska et al., 2024).

*Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies. Funding acknowledgements go at the end of the paper.

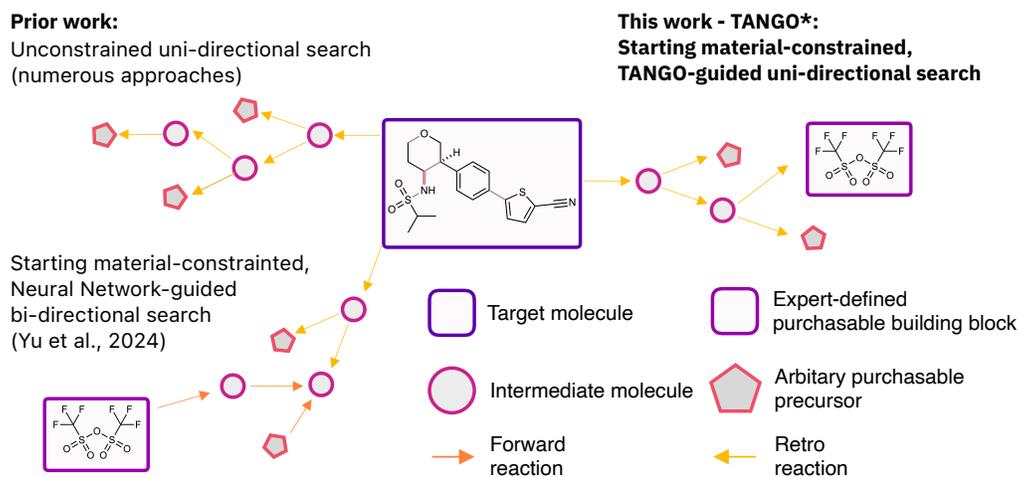


Figure 1: Comparison of existing constrained synthesis planning methods with Tango*

While designing constrained and steerable chemical synthesis is a daily practice in synthetic chemistry, it has received little attention in the CASP literature, with existing algorithms simply seeking to find any "valid" pathway to purchasable molecules.

Recently, several approaches for *starting material constrained* synthesis planning have been proposed with promising results (Johnson et al., 1992; Yu et al., 2022; 2024). Existing solutions either rely on rule-based approaches or require complex systems with several interacting parts. In this work, we show that a *general-purpose* and *data-driven* retrosynthesis system can be adapted to *starting material constrained* synthesis planning by the addition of a *computed* node cost function. Our contribution is as follows:

1. We use a computed node cost function, **TANimoto Group Overlap (TANGO)**, to guide the retrosynthetic search process toward enforced blocks. In this work, these blocks are limited to starting materials but could include key intermediates or molecular substructures.
2. We show that by integrating TANGO into an existing *general-purpose* search algorithm, we can tackle the constrained synthesis planning problem with results comparable to or superior to existing specialised methods.
3. We further show that TANGO can be a drop-in replacement for retrosynthetic cost networks and demonstrate that by doing so, the performance of the recently proposed bidirectional search algorithm (Yu et al., 2024) can be significantly improved.
4. We compare the outputs of existing retrosynthetic value functions with the outputs of the TANGO node cost function and present a plausible explanation for the improved performance over existing starting material constrained synthesis planning tools.

2 RELATED WORK

Computer-Assisted Synthesis Planning (CASP) tools typically formulate synthesis planning as a tree search, with each step corresponding to disconnecting a molecule into precursors through a retro" chemical reaction. Two primary approaches are used for selecting retro" reactions. Firstly, template-based methods extract chemical graph transformations from a corpus and train a neural network to select a transformation given an input (Segler & Waller, 2017; Segler et al., 2018). Template-free methods frame single-step retrosynthesis as a conditional language generation problem, with molecules encoded as SMILES strings or as a graph-edit prediction task (Weininger, 1988; Schwaller et al., 2020; Sacha et al., 2021; Zhong et al., 2023). Additionally, models that leverage graph features for direct generation have been developed (Tu & Coley, 2021; Chen et al., 2023). Significant focus has been placed on how to use single-step models in multi-step synthetic planning. Initial approaches used hand-curated rules, while more recent methods use neural-network guided graph exploration,

such as Monte Carlo Tree Search (MCTS) or AND-OR graph search methods (Pensak & Corey, 1977; Szymkuć et al., 2016; Segler et al., 2018; Genheden et al., 2020; Kishimoto et al.; Chen et al., 2020). A key development was driven by Chen et al. (2020), who proposed an A-star-like algorithm guided by a neural network that estimates the *cost to synthesise* a molecule from any arbitrary purchasable building block (Chen et al., 2020). More novel methods have utilised self-play and experience-based learning to improve navigation of the search space (Schreck et al., 2019; Kim et al., 2021; Hong et al., 2023; Liu et al., 2023). While single-step model performance continues to improve, this has not always been translated into the real-world performance of multi-step CASP systems (Thakkar et al., 2020; Schwaller et al., 2020; Maziarz et al., 2023; Torren-Peraire et al., 2024).

Constrained Single-Step Retrosynthetic Models. In recent years, there has been increased focus on introducing constraints into single-step retrosynthesis models with specific goals. Toniato *et al.* utilised reaction class tokens to steer the output of single-step retrosynthetic transformers towards specific reaction classes (Toniato et al., 2023). Following a similar approach, Thakkar *et al.* introduced "disconnection prompts" to guide single-step models to break specific bonds (Thakkar et al., 2023). In the multi-step planning domain, Westerlund *et al.* proposed a disconnection-aware transformer to encourage the breaking of bonds and allow the freezing of bonds during the search process, discarding any reaction that violates the frozen bond constraint (Westerlund et al., 2024). Interestingly, such bond constraints do not appear to impede the search process, indicating that simple, chemically informed rules can be powerful in data-driven retrosynthesis techniques.

Starting Material Constrained Synthesis Planning. Despite its potential use in waste valorisation and semi-synthesis, constrained synthesis planning has received limited attention in the literature. This approach imposes an additional constraint by focusing on the utilisation of specific starting materials. The LHASA program included such rules; however, they relied on expert-designed rules, limiting scalability (Johnson et al., 1992). GRASP utilised reinforcement learning to develop a *goal-driven* synthesis planning tool that can target either arbitrary products or specific starting materials (Yu et al., 2022). Recent state-of-the-art work proposed a bidirectional search algorithm, **Double-Ended Synthesis Planning (DESP)**, which uses both forward- and retro-expansion models, guided by a value network that estimates the cost of synthesising molecule *m2* specifically from molecule *m1* (Yu et al., 2024). DESP employs two search techniques, Frontier-to-Frontier (F2F) which compares all the nodes of each frontier with one another and Frontier-to-End (F2E) which compares the Retro* expansion frontier with the target starting material. Constrained synthesis planning has also emerged as a target in synthesisable molecular design. Guo et al. introduced a method for the *de novo* generation of synthesisable molecules using enforced building blocks in the synthesis pathway (Guo & Schwaller, 2024). To date, all starting material constrained synthesis planning tools have relied on specialised architectures, reinforcement learning, or expert-defined rules. In this work, we show instead that the problem can be approached with a simple cheminformatics calculation.

Algorithm 1: TANGO* Node Cost Function

Input :

MOL_{sm} ▷ Enforced Starting Materials
 $Node$ ▷ Molecule Node
 k ▷ Weight’s influence of TANGO to node cost
 c ▷ Scales TanSim vs. FMS

Function Tango*NodeCost ($Node, MOL_{sm}$) :

```

rewardnode ← 0
// Loop through all specified starting materials
foreach molsm ∈ MOLsm do
  // Compute reward for this starting material
  TanSim ← CalculateTanimotoSimilarity(node, molsm)
  FMS ← CalculateFMS(node, molsm)
  rewardsm ← TanSim · c + FMS · (1 - c)
  rewardnode ← max(rewardnode, rewardsm)
end
return k · (1 - rewardnode) + RetroCost

```

3 METHODS

TANGO* Overview

We adopt the traditional structure of the synthesis planning problem using a top-down, or retrosynthetic, search algorithm. Due to its strong performance and built-in compatibility with node cost-based guidance functions, we use Retro* (Chen et al., 2020) as the baseline search algorithm. The Retro* algorithm uses a best-first approach, selecting the lowest-cost node at any given iteration. In the original paper, a neural network that *estimates* the cost (number of chemical reactions) required to synthesise a given node from an arbitrary chemical building block is used to provide cost values for each node. To adapt the Retro* algorithm to the starting material constrained setting, we utilise a *computed* function, **TANGO**, to estimate the structural and chemical similarity of a node to a *specific* starting material. **TANGO** employs a weighted combination of Tanimoto Similarity and Fuzzy Matching Substructure (FMS) to compute molecular similarity. The Tango* cost is bounded by the interval $([0, 1])$, while the Retro* cost is theoretically bounded by $([0, \infty))$ and practically varies around the range $([0, 10])$. To balance the influence of these mechanisms in the search dynamics, a scalar hyperparameter (k) is used to upweight the Tango* cost. The two cost functions are additively combined; further details are provided in Algorithm 1 and Appendix A.

4 EXPERIMENTS

Tango* is built on the Retro* algorithm, using a reference implementation provided by DESP (Yu et al., 2024). For our estimation of node cost, we combine the original Retro* Value function (Chen et al., 2020) with an adaptation of the TANGO reward introduced by Guo & Schwaller (2024).

Our experiments are structured to answer the following questions: **(1)** Can a non-neural computed node cost function be used to adapt general-purpose synthesis planning tools to the constrained setting? **(2)** Can such a system outperform existing specialised models for starting material-constrained planning? **(3)** Can the cost function additionally provide improvements to existing bidirectional search methods? **(4)** As the TANGO function is empirically *computed* as opposed to *estimated*, does TANGO generalise from simple to harder datasets more effectively?

4.1 EXPERIMENTAL SETUP

Datasets. To evaluate our system’s performance, we utilise the common **USPTO-190** dataset introduced by Chen et al. (2020), which is a set of 190 challenging target molecules extracted from USPTO-Full. Additionally, we use the datasets introduced by Yu et al. (2024), namely **Pistachio Reachable** and **Pistachio Hard**. The sets of *target*, *starting material* pairs are extracted for a set of commercial building blocks, we use canonical SMILES strings provided in the set of 23 million molecules from eMolecules used by Chen et al. (2020) and Yu et al. (2024).

Machine learning models. To avoid variance due to subtle differences in data pre-processing techniques and to ensure a meaningful comparison, we use the Retro* value network and single-step retrosynthesis model provided by Yu et al. (Yu et al., 2024; Maziarz et al., 2023).

Hyperparameter Optimisation. We use a hyperparameter k to balance the starting material guidance of TANGO with the general guidance of the Retro* Value Network. To evaluate the ability of our method to generalise from simpler to more complex molecules, we choose the Pistachio Reachable dataset for hyperparameter tuning. We find a value of $k = 25$ optimises both Solve Rate and Average Number of expansions. We employ an additional parameter, c , to specify the ratio of FMS to Tanimoto Similarity, with c defining the FMS weight. Through empirical testing, we determine the optimal value to be $c = 0.3$. In the results section, we will refer to Tango with $c = 0.0$ as Tango(1, 0) and Tango with $c = 0.3$ as Tango(0.7, 0.3).

4.2 TANGO* SOLVE RATE AND ALGORITHM EFFICIENCY

While there remains a lack of a clearly agreed-upon ‘gold standard’ for reporting the quality of synthetic routes generated by CASP systems, a few metrics are commonly used. The **Solve Rate** indicates the system’s ability to find viable solutions, while the **Average Route Length** (in terms of the number of reactions) indicates route quality. In addition to reporting metrics regarding the outputs

Table 1: Summary comparison between baseline methods and Tango* across the three benchmarks. Baseline results for Retro*, GRASP, Retro* + D and the DESP methods are taken from Yu et al. (2024). Solve rate is the fraction of (*target, starting material*) pairs solved within the expansion limit. Tango DESP methods use a (1, 0) weighting of Tanimoto Similarity to FMS. Best overall results are in **bold** and best uni-directional results are underlined

| Algorithm | USPTO-190 | | | | Pistachio Reachable | | | | Pistachio Hard | | | |
|------------------|---------------------------|-------------|-------------|------------------------|---------------------------|-------------|-------------|------------------------|---------------------------|-------------|-------------|------------------------|
| | Solve Rate (%) \uparrow | | | \bar{N} \downarrow | Solve Rate (%) \uparrow | | | \bar{N} \downarrow | Solve Rate (%) \uparrow | | | \bar{N} \downarrow |
| Expansion Budget | 100 | 300 | 500 | | 50 | 100 | 300 | | 100 | 300 | 500 | |
| Random | 4.2 | 4.7 | 4.7 | 479 | 16.0 | 26.7 | 40.7 | 325 | 6.0 | 12.0 | 13.0 | 452 |
| BFS | 12.1 | 20.0 | 24.2 | 413 | 48.7 | 57.3 | 74.0 | 169 | 16.0 | 26.0 | 29.0 | 390 |
| MCTS | 20.5 | 32.1 | 35.3 | 364 | 52.0 | 72.7 | 85.3 | 111 | 27.0 | 31.0 | 32.0 | 361 |
| Retro* | 25.8 | 33.2 | 35.8 | 351 | 70.7 | 78.0 | 92.7 | 73 | 32.0 | 35.0 | 37.0 | 342 |
| GRASP | 15.3 | 21.1 | 23.7 | 410 | 46.7 | 51.3 | 66.7 | 198 | 14.0 | 22.0 | 29.0 | 402 |
| Retro*+D | 27.4 | 32.6 | 37.4 | 348 | 77.3 | 87.3 | 96.0 | 49 | 31.0 | 40.0 | 42.0 | 323 |
| DESP-F2E | 30.0 | 35.3 | 39.5 | 340 | 84.0 | 90.0 | 96.0 | 41 | 35.0 | 44.0 | 50.0 | 300 |
| DESP-F2F | 29.5 | 34.2 | 39.5 | 336 | 84.5 | 88.9 | 97.3 | 38 | 39.0 | 45.0 | 48.0 | 293 |
| Ours | | | | | | | | | | | | |
| Tango(1, 0)* | 36.3 | <u>41.1</u> | <u>42.6</u> | <u>313</u> | 84.5 | 90.6 | 97.3 | 32 | <u>40.0</u> | <u>45.0</u> | <u>47.0</u> | <u>290</u> |
| Tango(0.7, 0.3)* | 35.7 | 40.5 | 42.1 | 316 | <u>86.7</u> | <u>94.0</u> | <u>98.0</u> | <u>29</u> | 39.0 | 44.0 | 46.0 | 295 |
| Tango-F2E | 33.1 | 40.0 | 41.5 | 317 | 88.7 | 92.0 | 98.7 | 27 | 39.0 | 45.0 | 49.0 | 290 |
| Tango-F2F | 33.2 | 45.3 | 53.7 | 291 | 91.3 | 95.3 | 99.3 | 18 | 47.0 | 59.0 | 63.0 | 231 |

Table 2: Inference time and Mean Solved Route Length for the evaluated methods. Route length comparisons are made on the routes solved by all methods.

| Algorithm | USPTO-190 | | Pistachio Reachable | | Pistachio Hard | |
|--------------|-----------------------------|------------------------|------------------------------|------------------------|-----------------------------|------------------------|
| | Route Length (61 Routes) | Wall Clock Time (s) | Route Length (114 Routes) | Wall Clock Time (s) | Route Length (36 Routes) | Wall Clock Time (s) |
| Retro* | 5.30 | 58.1 | 4.64 | 10.2 | 4.67 | 56.3 |
| Retro* + D | 5.56 | 64.1 | 4.67 | 8.3 | 4.67 | 55.2 |
| DESP-F2E | 5.13 | 66.5 | 4.51 | 8.6 | 4.56 | 56.3 |
| DESP-F2F | 5.51 | 109.4 | 4.46 | 8.2 | 4.44 | 61.8 |
| Tango(1, 0)* | 5.06 | 55.8 | 4.56 | 5.8 | 4.67 | 47.5 |
| Tango-F2E | 4.44 | 75.3 | 4.24 | 6.5 | 4.29 | 54.2 |
| Tango-F2F | 5.06 | 146.4 | 4.04 | 4.9 | 4.40 | 72.8 |

of Tango*, we aim to evaluate how efficient Tango* is at navigating the retrosynthetic search space and assess the effect of computational overhead on the system. To this end, we report the **Average Number of Expansions** (\bar{N}) and **Wall Clock Time**.

The primary results are displayed in Table 1. Tango(1, 0)* demonstrates improvements in the starting material-constrained setting and consistently outperforms the neural network-enhanced Retro* (referred to as Retro* + D) across all benchmarks and expansion limits while displaying greater algorithmic efficiency, as measured by the average number of expansions \bar{N} . In addition, Tango(1, 0)* achieves higher or comparable solve rates to both DESP methods across all three datasets, doing so with a strictly lower average number of expansion calls, clearly demonstrating the power of the TANGO reward to navigate the retrosynthetic action space. We find that the value of k , optimised on Pistachio Reachable, shows strong generalisation performance to the more challenging datasets, with Tango* being the best-performing method across all iteration levels on USPTO-190.

We perform an ablation of the FMS : Tanimoto Similarity weighting in the **TANGO** cost function, which we refer to as Tango(0.7, 0.3). We find that although the incorporation of FMS into the cost function improves the solve rate and reduces expansion calls for Pistachio Reachable, such results do not carry over to the more challenging datasets. We hypothesise that Tanimoto Similarity offers greater granularity for guidance than FMS, enabling higher performance on more complex datasets.

To examine the general applicability of the TANGO cost in guiding various search algorithms, we explore its integration into the recently proposed bidirectional search methods, DESP-F2F and

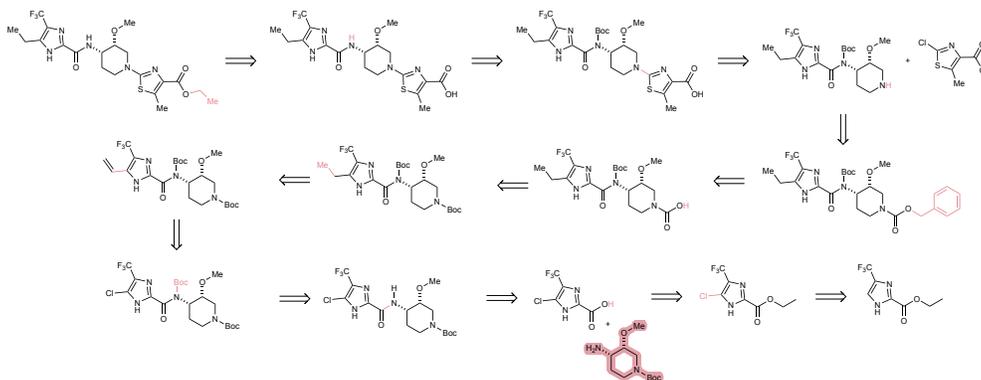


Figure 2: Here we demonstrate a meaningful 12-step route generated by our method on a (*target, starting material*) pair not solved by the best performing DESP (Yu et al., 2024) method. Constrained starting material highlighted in red; bonds/atoms disconnected shown in red.

DESP-F2E (Yu et al., 2024). This integration is achieved by replacing the pairwise synthetic distance network, D , with the TANGO cost function. The hyperparameters k and c are set to the same values as Tango (1,0). Our findings show that the addition of TANGO reward generally leads to a substantial increase in the solve rate for both DESP methods, while also reducing the average number of expansions and route length. Particularly noteworthy is the impressive performance of Tango-F2F at high expansion budgets, where it achieves a 99.3% solve rate on Pistachio Reachable and improves accuracy by approximately 25% compared to the next best method on the more challenging UPSTO-190 and Pistachio Hard datasets. We note that as the added bidirectional search of DESP outperforms Retro* + D , TANGO-DESP methods should be expected to outperform Tango*.

4.3 WALL CLOCK TIME AND ROUTE LENGTH ABLATIONS

As both DESP and Tango* introduce computational overhead that may decouple the number of expansion calls from the computational resources required, we report the wall clock time. Tango* consistently achieves a lower wall clock time than alternative starting material-constrained methods. Finally, we investigate the average number of reactions per solved route for each method. Tango* achieves shorter route lengths than all existing methods on USPTO-190, but only matches existing methods on other datasets. The strongest results in terms of route length come from the combination of TANGO with bidirectional search methods F2E and F2F, with one of them displaying the shortest routes for all of the datasets. This result is revealing; existing methods use a neural network directly trained to predict synthetic distance, and yet it fails to provide significantly stronger guidance towards shorter synthetic routes than a simply molecular similarity measure. This leads to the obvious question, just how effective are such neural networks at estimating the synthetic distance of a node, and how reliable is this estimation at test time?

4.4 WHY DOES TANGO* WORK?

Despite access to a *starting material constrained* node cost function (one with access to information from ground truth routes at test time), Retro* + D does not show a substantial increase in performance compared to Retro*. In contrast, Tango*'s incorporation of a privileged node cost function provides significant performance improvements. We hypothesize that, as a computed cost function, the TANGO cost function should be relatively invariant to the molecular inputs and maintain strong performance at test time. Let $V : M \rightarrow \mathbb{R}$ be a node cost function where M is the space of molecules. For node cost to effectively guide retrosynthetic search the function should *generally* satisfy:

$$V(m_i) > V(m_{i+1}) \quad \forall i \in 1, \dots, n-1 \quad (1)$$

where m_1, \dots, m_n represents molecules along a synthetic path from root to starting material.

A potential caveat to the above constraint might be "tactical combinations", the synthetic complexity of a molecule temporarily increases during a retrosynthesis, to allow a major complexity-reducing reaction Gajewska et al. (2020). We believe that retrosynthetic value functions should continue to

show synthetic distance estimate decreases even for the complexity-increasing step. We also expect a stronger monotonic decrease for routes that are *solved* by a method, compared to routes *not solved*, as this indicates that the node cost function can accurately prioritise nodes in the search tree. To systematically evaluate these hypotheses about TANGO’s effectiveness and empirically assess the relative strength of different guidance functions, we analyse their behavior on ground truth synthetic routes in the test set. Using the USPTO-190 dataset, we extract the linear synthetic path from root molecule m_r to expert-defined starting material m_s . We define synthetic distance $d(m_1, m_2)$ as the minimum number of reactions required to synthesize m_1 from m_2 . For each molecule m_i in this path, we calculate:

- $d(m_i, m_s)$: Ground truth synthetic distance
- $D(m_i, m_s)$: Neural network estimation of synthetic distance
- $T(m_i, m_s)$: TANGO cost molecular similarity

To isolate the cost function, we fix the search algorithm, in this case focusing on the Retro* algorithm with either TANGO or D as a starting material-guided cost function. We then take the 4 sets of routes that are solved and not solved by Tango* and Retro* + D . We plot the corresponding starting material constrained cost function, TANGO and D respectively, as costs for each node in the ground truth synthetic routes.

We show the results of this experiment in Figure 3. The neural network estimated synthetic distance **(b)** and **(e)** displays an unexpected bimodal distribution, with peaks at low and high estimates. It displays consistently high absolute error and is unable to provide a granular, monotonically decreasing estimate of synthetic distance. For routes that Retro* + D solves, the mean synthetic distance estimate begins to decrease when the ground truth distance is under two but still displays substantial error margins. This is exacerbated for routes that Retro* does not solve, with the synthetic distance estimation consistently varying around 10 if the ground truth distance is greater than 1. We note this fits with the training strategy of D described in (Yu et al., 2024), which augments the training set with synthetic "negative samples" of (*target, startingmaterial*) pairs. These samples are generated by selecting two molecules that are disconnected in the directed graphs formed by linking reactants and products in USPTO, and are assigned a fixed "distance" value of 10. While such a method should be robust in a *dense* dataset where all theoretically synthesisable molecules are interconnected, USPTO is notably *sparse*. Consequently, many of these "disconnected" samples may, in fact, be feasibly synthesisable. Furthermore, the constraint enforcing a Tanimoto Similarity of less than 0.7 may be overly permissive. We observe that the majority of molecules within the same synthetic route typically exhibit Tanimoto similarities far lower than 0.7, suggesting that this sampling approach may introduce significant noise into the training process.

In comparison, the computed TANGO cost function **3(a)** and **3(d)** exhibits a clear and granular monotonic decrease across routes solved by Tango*. On routes not solved by Tango*, the monotonic decrease is less pronounced but still present. We hypothesize that this strong granularity and monotonicity enables TANGO cost-guided search algorithms to achieve substantially improved solve rates compared to neurally guided methods. Based on this observation, we investigated whether similar findings apply to *non-privileged* (non-goal oriented) neural guidance functions, such as the original Retro* cost function. Prior work has shown that omission of the Retro* cost function has limited effects on retrosynthetic search performance, with both Chen et al. (2020) and Maziarz et al. (2023) finding that setting the value to a constant resulted in only minor performance changes on USPTO-190. The results for solved and unsolved routes are shown in Figure 3 **(c)** and **(f)**, respectively. In contrast to the plots for the D cost function, which consistently *overestimates* cost, the Retro* cost function consistently *underestimates* it. While it displays clear monotonicity on solved routes, this pattern does not extend to unsolved routes, where the cost function hovers around 1-2 before splitting into a bimodal distribution with *lower error* at ground truth distances greater than 7. This is likely a function of the power-law like distribution of reaction sequence lengths in USPTO, with the majority of synthetic routes consisting of only a single step. These results indicate that value networks trained exclusively on positive samples using simple training methods such as MSE and Consistency loss perform substantially better at estimating synthetic distances.

4.5 CASE STUDY: SYNTHESIS OF USEFUL COMPOUNDS FROM RENEWABLE/WASTE FEEDSTOCK

A key aim of starting-material-constrained synthesis planning is to enable the discovery of synthetic pathways to useful compounds from renewable or waste feedstocks. Previous evaluations relied

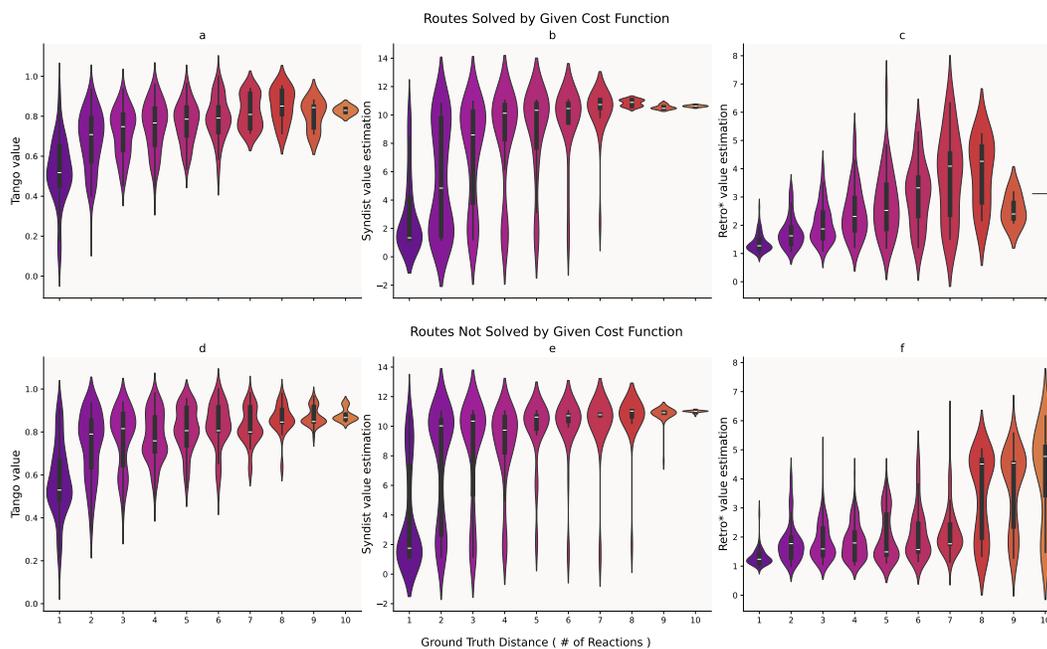


Figure 3: A comparison of node cost estimates for USPTO-190 routes solved and not solved by Retro* search using the corresponding cost function. **(a)** Tango Cost for routes solved by Tango*, **(b)** SynDist Cost for routes solved by Retro* + D, **(c)** Retro* cost function estimates for routes solved by Retro*, **(d)** Tango Cost for routes not solved by Tango*, **(e)** SynDist cost for routes not solved by Retro* + D and **(f)** Retro* cost for routes not solved by Retro*

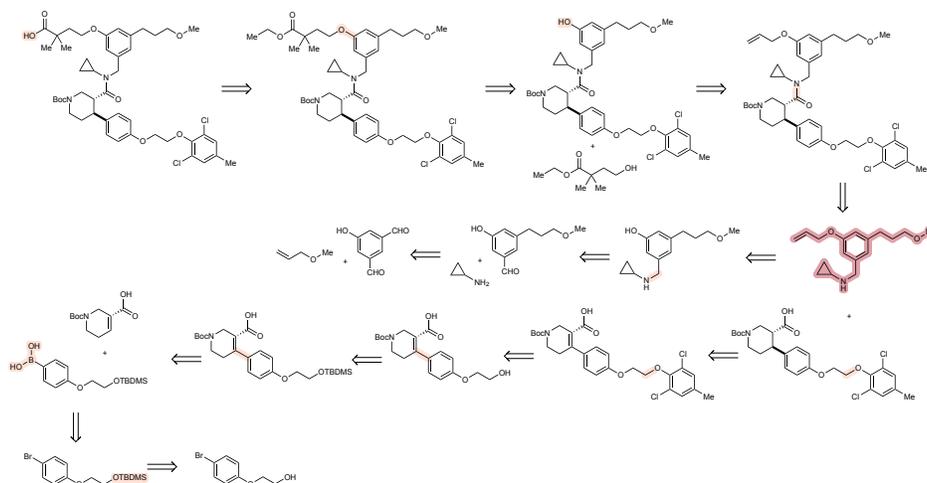


Figure 4: Here we demonstrate a feasible 10-step route generated by Tango-DESP-F2F on a (*target, starting material*) pair not solved by the neural guided DESP-F2F method. Constrained starting material is highlighted in red; bonds/atoms disconnected are shown in red.

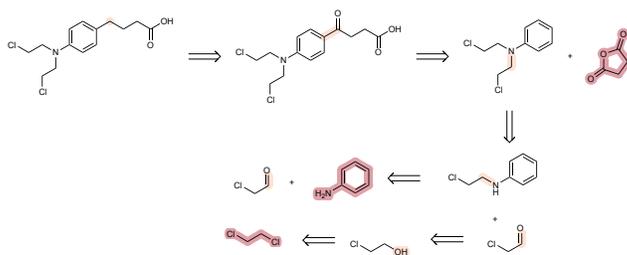


Figure 5: Here we show a feasible synthesis route to the chemotherapy drug, Chlorambucil, a WHO essential medicine, synthesised entirely from renewable or industrial waste feedstocks.

on a database of 17 million chemical building blocks from the eMolecules database. We aim to demonstrate the effectiveness of our method, Tango*, in finding synthesis pathways to useful small molecules starting *exclusively* from renewable or waste feedstocks. For renewable building blocks, we use a set of 146 small molecules previously curated by Wołos et al. (2022). For useful compounds, we extract a set of 110 small molecules from a curation of the WHO List of Essential Medicines, previously developed by Gao et al. (2020). We conduct the search using the Tango(1, 0) set of hyperparameters previously described, but set the expansion budget to 1,000 model calls. In Figure 5, we present a strong route, discovered by Tango* but not Retro*, to the chemotherapy drug Chlorambucil, starting exclusively from renewable starting materials. We find all proposed reactions and the complete synthesis are directly reported in the literature (Yang et al., 2020; Gangjee et al., 2009)

5 CONCLUSION

In this work, we introduce Tango*, a simple adaptation of the Retro* algorithm to the starting material-constrained setting without any model retraining. We demonstrate that our TANGO guided search method strictly outperforms the similar neural network-guided Retro* + *D*. Despite relying on single-ended search, Tango* either outperforms or matches the performance of specialised DESP models and search algorithms, providing routes that satisfy the specified goal for a greater number of compounds. Application of the TANGO node cost function to the DESP methods also yields substantial improvements, particularly to the F2F method, which achieves the strongest solve rate performance of all investigated systems. It proposes routes with a comparable length and does so with a lower number of expansion calls and reduced wall clock time.

We show that existing neural node cost functions fail to provide a granular and monotonic decrease in node cost throughout a retrosynthesis pathway, particularly struggling on more challenging routes. In contrast, the computed Tango* cost function displays better monotonicity and granularity on both solved and unsolved routes. This work indicates that there may be substantial room for improvement in developing novel guidance functions for retrosynthesis tools.

We anticipate that future developments in similar methods will unlock synthesis planning tools with diverse and flexible structure constraints, allowing expert chemists to specify key intermediates or predefined substructure goals at any position in the synthetic route.

Code and Data availability The anonymised code used to produce these results is available here <https://figshare.com/s/5425c38da7a686d1f878>

Acknowledgements This research was funded by the Swiss National Science Foundation (SNSF) [214915]. Additional funding was provided by NCCR Catalysis (grant number 225147), a National Centre of Competence in Research funded by the Swiss National Science Foundation.

REFERENCES

- Zachary G Brill, Matthew L Condakes, Chi P Ting, and Thomas J Maimone. Navigating the chiral pool in the total synthesis of complex terpene natural products. *Chemical reviews*, 117(18): 11753–11795, 2017.
- Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: Learning retrosynthetic planning with neural guided A* search. In *International Conference on Machine Learning*, pp. 1608–1616. PMLR, 2020.
- Ziqi Chen, Oluwatosin R Ayinde, James R Fuchs, Huan Sun, and Xia Ning. G 2 retro as a two-step graph generative models for retrosynthesis prediction. *Communications Chemistry*, 6(1):102, 2023.
- Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Computer-assisted retrosynthesis based on molecular similarity. *ACS Cent. Sci.*, 3:1237–1245, 2017.
- E. J. Corey and Xue-Min Chelg. *The Logic of Chemical Synthesis*. John Wiley & Sons, 1995.
- Elias James Corey. General methods for the construction of complex molecules. *Pure and Applied chemistry*, 14:19–38, 1967.
- Elias James Corey and W Todd Wipke. Computer-assisted design of complex organic syntheses: Pathways for molecular synthesis can be devised with a computer and equipment for graphical communication. *Science*, 166(3902):178–192, 1969.
- Ewa P. Gajewska, Sara Szymkuć, Piotr Dittwald, Michał Startek, Oskar Popik, Jacek Mlynarski, and Bartosz A. Grzybowski. Algorithmic Discovery of Tactical Combinations for Advanced Organic Syntheses. *Chem*, 6(1):280–293, 2020. ISSN 2451-9294.
- Aleem Gangjee, Ying Wang, Sherry F. Queener, and Roy L. Kisliuk. Synthesis of 2,6-diamino-5-[(2-substituted phenylamino)ethyl]pyrimidin-4(3h)-one as inhibitors of folate metabolizing enzymes. *Journal of Heterocyclic Chemistry*, 43, March 2009.
- Hanyu Gao, Connor W Coley, Thomas J Struble, Linyan Li, Yujie Qian, William H Green, and Klavs F Jensen. Combining retrosynthesis and mixed-integer optimization for minimizing the chemical inventory needed to realize a who essential medicines list. *Reaction Chemistry & Engineering*, 5: 367–376, 2020.
- Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning. *J. Cheminf.*, 12:1–9, 2020.
- Jeff Guo and Philippe Schwaller. It takes two to tango: Directly optimizing for constrained synthesizability in generative molecular design. *arXiv preprint arXiv:2410.11527*, 2024.
- Siqi Hong, Hankz Hankui Zhuo, Kebing Jin, Guang Shao, and Zhanwen Zhou. Retrosynthetic planning with experience-guided monte carlo tree search. *Communications Chemistry*, 6(1):120, 2023.
- A Peter Johnson, Chris Marshall, and Philip N Judson. Starting material oriented retrosynthetic analysis in the lhasa program. 1. general description. *Journal of chemical information and computer sciences*, 32(5):411–417, 1992.
- Junsu Kim, Sungsoo Ahn, Hankook Lee, and Jinwoo Shin. Self-improved retrosynthetic planning. *arXiv preprint arXiv:2106.04880*, 2021.
- Akihiro Kishimoto, Beat Buesser, Bei Chen, and Adi Botea. Depth-first proof-number search with heuristic edge cost and application to chemical synthesis planning.
- Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS Cent. Sci.*, 3:1103–1113, 2017.

- Guoqing Liu, Di Xue, Shufang Xie, Yingce Xia, Austin Tripp, Krzysztof Maziarz, Marwin Segler, Tao Qin, Zongzhang Zhang, and Tie-Yan Liu. Retrosynthetic planning with dual value networks. In *International Conference on Machine Learning*, pp. 22266–22276. PMLR, 2023.
- Lauren M Lopez, Quan Zhang, Orion Dollar, Jim Pfaendtner, Brent H Shanks, and Linda J Broadbelt. Application of automated network generation for retrosynthetic planning of potential corrosion inhibitors. *Molecular Systems Design & Engineering*, 9(4):352–371, 2024.
- Krzysztof Maziarz, Austin Tripp, Guoqing Liu, Megan Stanley, Shufang Xie, Piotr Gaiński, Philipp Seidl, and Marwin Segler. Re-evaluating retrosynthesis algorithms with syntheseus. *arXiv preprint arXiv:2310.19796*, 2023.
- Iwao Ojima, Ivan Habus, Mangzhu Zhao, Martine Zucco, Young Hoon Park, Chung Ming Sun, and Thierry Brigaud. New and efficient approaches to the semisynthesis of taxol and its c-13 side chain analogs by means of β -lactam synthon method. *Tetrahedron*, 48(34):6985–7012, 1992.
- David A Pensak and Elias James Corey. Lhasa—logic and heuristics applied to synthetic analysis. ACS Publications, 1977.
- Mikolaj Sacha, Mikolaj Blaz, Piotr Byrski, Pawel Dabrowski-Tumanski, Mikolaj Chrominski, Rafal Loska, Pawel Wlodarczyk-Pruszynski, and Stanislaw Jastrzebski. Molecule edit graph attention network: Modeling chemical reactions as sequences of graph edits. *J. Chem. Inf. Model.*, 61: 3273–3284, 2021. PMID: 34251814.
- Nadine Schneider, Daniel M Lowe, Roger A Sayle, Michael A Tarselli, and Gregory A Landrum. Big data from pharmaceutical patents: a computational analysis of medicinal chemists’ bread and butter. *Journal of medicinal chemistry*, 59(9):4385–4402, 2016.
- John S Schreck, Connor W Coley, and Kyle JM Bishop. Learning retrosynthetic planning through simulated experience. *ACS central science*, 5(6):970–981, 2019.
- 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. *Chem. Sci.*, 11: 3316–3325, 2020.
- Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chem. Eur. J.*, 23:5966–5971, 2017.
- Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018.
- Sara Szymkuć, Ewa P Gajewska, Tomasz Klucznik, Karol Molga, Piotr Dittwald, Michał Startek, Michał Bajczyk, and Bartosz A Grzybowski. Computer-assisted synthetic planning: the end of the beginning. *Angew. Chem. - Int. Ed.*, 55:5904–5937, 2016.
- Amol Thakkar, Thierry Kogej, Jean-Louis Reymond, Ola Engkvist, and Esben Jannik Bjerrum. Datasets and their influence on the development of computer assisted synthesis planning tools in the pharmaceutical domain. *Chem. Sci.*, 11:154–168, 2020.
- Amol Thakkar, Alain C Vaucher, Andrea Byekwaso, Philippe Schwaller, Alessandra Toniato, and Teodoro Laino. Unbiasing retrosynthesis language models with disconnection prompts. *ACS Central Science*, 9(7):1488–1498, 2023.
- Alessandra Toniato, Alain C Vaucher, Philippe Schwaller, and Teodoro Laino. Enhancing diversity in language based models for single-step retrosynthesis. *Digital Discovery*, 2(2):489–501, 2023.
- Paula Torren-Peraire, Alan Kai Hassen, Samuel Genheden, Jonas Verhoeven, Djork-Arné Clevert, Mike Preuss, and Igor V. Tetko. Models matter: the impact of single-step retrosynthesis on synthesis planning. *Digital Discovery*, 3:558–572, 2024.
- Zhengkai Tu and Connor W Coley. Permutation invariant graph-to-sequence model for template-free retrosynthesis and reaction prediction. *arXiv preprint arXiv:2110.09681*, 2021.

David Weininger. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.*, 28:31–36, 1988.

Annie M Westerlund, Lakshidaa Saigiridharan, and Samuel Genheden. Constrained synthesis planning with disconnection-aware transformer and multi-objective search. 2024.

Agnieszka Wołos, Dominik Koszelewski, Rafał Roszak, Sara Szymkuć, Martyna Moskal, Ryszard Ostaszewski, Brenden T. Herrera, Josef M. Maier, Gordon Brezicki, Jonathon Samuel, Justin A. M. Lummiss, D. Tyler McQuade, Luke Rogers, and Bartosz A. Grzybowski. Computer-designed repurposing of chemical wastes into drugs. *Nature*, 604(7907):668–676, 2022.

Zhiheng Yang, Xueya Zhu, Shiyi Yang, Weiyang Cheng, Xiaojian Zhang, and Zhanhui Yang. Iridium-catalysed reductive deoxygenation of ketones with formic acid as traceless hydride donor. *Advanced Synthesis & Catalysis*, 362(23):5496–5505, 2020.

Kevin Yu, Jihye Roh, Ziang Li, Wenhao Gao, Runzhong Wang, and Connor W Coley. Double-ended synthesis planning with goal-constrained bidirectional search. *arXiv preprint arXiv:2407.06334*, 2024.

Yemin Yu, Ying Wei, Kun Kuang, Zhengxing Huang, Huaxiu Yao, and Fei Wu. Grasp: Navigating retrosynthetic planning with goal-driven policy. *Advances in Neural Information Processing Systems*, 35:10257–10268, 2022.

Anna Źadło-Dobrowolska, Karol Molga, Olga O Kolodiazna, Sara Szymkuć, Martyna Moskal, Rafał Roszak, and Bartosz A Grzybowski. Computational synthesis design for controlled degradation and revalorization. *Nature Synthesis*, pp. 1–12, 2024.

Weihe Zhong, Ziduo Yang, and Calvin Yu-Chian Chen. Retrosynthesis prediction using an end-to-end graph generative architecture for molecular graph editing. *Nature Communications*, 14(1):3009, 2023.

A APPENDIX

A.1: Compute Details

Wall clock time comparisons were implemented on a GPU-enabled workstation with the following specifications.

- CPU: 12-core AMD Ryzen 9 7900X
- RAM Memory : 64 GB
- GPU : NVIDIA A6000 48 GB

A.2: Retro* Algorithm

We note this algorithm description is largely taken from the DESP paper Yu et al. (2024).

Retro* defines the following quantities:

- V_m : For a molecule m , V_m is an unconditional estimate of the minimum cost required to synthesise m . This estimate is provided by a neural network.
- $\text{rn}(m|G)$: For a molecule m in the search graph G , the "reaction number" $\text{rn}(m|G)$ represents the estimated minimal cost to synthesise m .
- $V_t(m|G)$: For a molecule m in the search graph G with the target molecule p^* , $V_t(m|G)$ denotes the estimated minimal cost to synthesise p^* starting from m .

Retro* operates through iterative phases of selection, expansion, and update. We follow the DESP implementation of Retro* as follows:

Selection: Choose the molecule from the set of frontier nodes F that minimises the expected cost of synthesising the target p^* given the current search graph G :

$$m_{\text{select}} = \arg \min_{m \in F} V_t(m|G) \quad (2)$$

Expansion: As detailed in Algorithm 2, apply a one-step retrosynthesis model to the selected node m_{select} , and add the resulting reactions and precursor molecules to the search graph G . Initialise each new molecule node with:

$$\text{rn}(m|G) \leftarrow V_m$$

Update:

First, propagate the reaction number values upward to ancestor nodes. For a reaction node R , update its reaction number as the sum of its childrens’ reaction numbers plus the cost of the reaction $c(R)$:

$$\text{rn}(R|G) \leftarrow c(R) + \sum_{m \in \text{ch}(R)} \text{rn}(m|G) \quad (3)$$

For a molecule node m , update its reaction number to be the minimum reaction number among its child reactions:

$$\text{rn}(m|G) \leftarrow \min_{R \in \text{ch}(m)} \text{rn}(R|G) \quad (4)$$

Next, propagate the values of $V_t(m|G)$ downward to descendant nodes. Starting from the target molecule p^* :

$$V_t(p^*|G) \leftarrow \text{rn}(p^*|G) \quad (5)$$

For subsequent reaction nodes R , update the value as:

$$V_t(R|G) \leftarrow \text{rn}(R|G) - \text{rn}(\text{pr}(R)|G) + V_t(\text{pr}(R)|G) \quad (6)$$

Finally, for molecule nodes m that are not the target p^* :

$$V_t(m|G) \leftarrow \min_{R \in \text{pr}(m)} V_t(R|G) \quad (7)$$

Here, $\text{ch}(R)$ denotes the set of child molecules of reaction R , and $\text{ch}(m)$ represents the set of child reactions of molecule m . Similarly, $\text{pr}(R)$ denotes the parent molecule of reaction R , and $\text{pr}(m)$ represents the set of parent reactions of molecule m .

This implementation ensures that at each iteration, the algorithm selects the most promising node to expand based on the estimated cost, propagates cost updates throughout the search graph, and efficiently guides the search towards the most cost-effective synthesis pathways.

A.3: Dataset Details

We provide details on the used datasets in Table 3. The Table is taken directly from the original DESP paper (Yu et al., 2024).

A.4: TANGO Hyperparameter Screening

Here we provide a hyperparameter screen of TANGO weight and Tanimoto weight across the 3 Tango augmented methods, **Tango***, **TANGO DESP-F2E** and **TANGO DESP-F2F**. Due to computational resource limitations, this screen was conducted exclusively on Pistachio Reachable with a small expansion budget of 50. In general, a higher Tango Weight (> 15) tends to lead to a higher solve rate, consistent across all methods.

A Tanimoto weight of 0.75 performs strongly for the **TANGO DESP-F2F** and **Tango*** methods on the Pistachio Reachable test set, consistent with previous results.

Table 3: Benchmark dataset summary. Avg. In-Dist. % is the mean percentage of reactions in each route within the top 50 suggestions of the retro model. Unique Rxn.% is the ratio of deduplicated reactions to total reactions across all routes. Avg. # Rxns. is the mean number of reactions in each route, and Avg. Depth is the mean number of reactions in the longest path of each route.

| Dataset | # Routes | Avg. In-Dist. % | Unique Rxn. % | Avg. # Rxns. | Avg. Depth |
|---------------------|----------|-----------------|---------------|--------------|------------|
| USPTO-190 | 190 | 65.6 | 50.5 | 6.7 | 6.0 |
| Pistachio Reachable | 150 | 100 | 86.1 | 5.5 | 5.4 |
| Pistachio Hard | 100 | 59.9 | 95.2 | 7.5 | 7.2 |

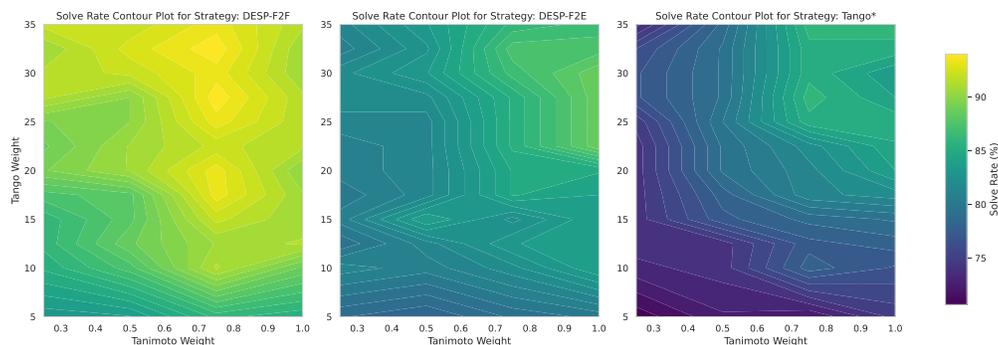


Figure 6: A hyper-parameter screen conducted on Pistachio Reachable with an expansion limit of 50.

A.5: Renewable synthesis of WHO Essential Medicines datasets

The sources of small molecule WHO Essential Medicines and renewable and waste compound feeds Gao et al. (2020); Wołos et al. (2022) are not machine-readable. We used a vision-enabled large language model (GPT-4o) to parse compound names from the Supplementary Information sections of the aforementioned papers.

A.6: Value function failure points

In Figure 7, we analyze discrepancies between model-predicted number of synthesis steps (Retro* and SynthDist) and ground truth values across four synthetic transformations. In these examples, models overestimate synthetic complexity compared to ground truth, likely due to their focus on structural features rather than strategic synthetic planning. This is evident in: (a) ring-opening of tetrahydropyran introducing rotatable bonds of an alkyl chain, (b) TBDMS protection adding heavy atoms, (c) oxidation of primary alcohols changing atom connectivity and adding heavy atoms, and (d) lactone formation creating additional ring complexity. These cases demonstrate how the models may misinterpret strategic intermediates with increased synthetic complexity as being further away from commercially available building blocks while in reality, these intermediates help disconnections over multiple steps, suggesting potential areas for improvement in retrosynthetic prediction algorithms.

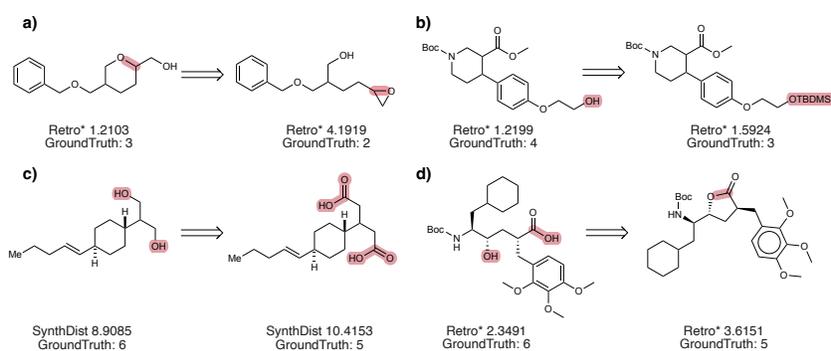


Figure 7: Failure points of value networks (Retro* and SynthDist) compared with ground truth synthetic distance values. Bonds and atoms that are being modified are highlighted.