Planning chemical syntheses with deep neural networks and symbolic AI

Marwin H. S. Segler, Mike Preuss & Mark P. Waller
(*Nature, March 2018)*

Presented by Edward Wu, Aryan Saha

Outline

- Context: What is retrosynthesis
- 3N-MCTS
  - 3N
  - MCTS
- Results
- Discussion

Context: What is retrosynthesis

- 1990 Nobel Prize, E. J. Corey
- Working backwards to build a molecule forward
- Path considered solved iff every single molecule in its broken down into is in

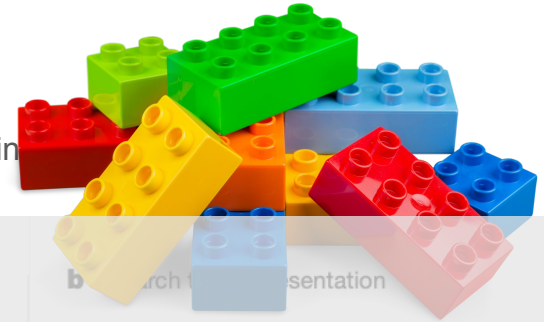

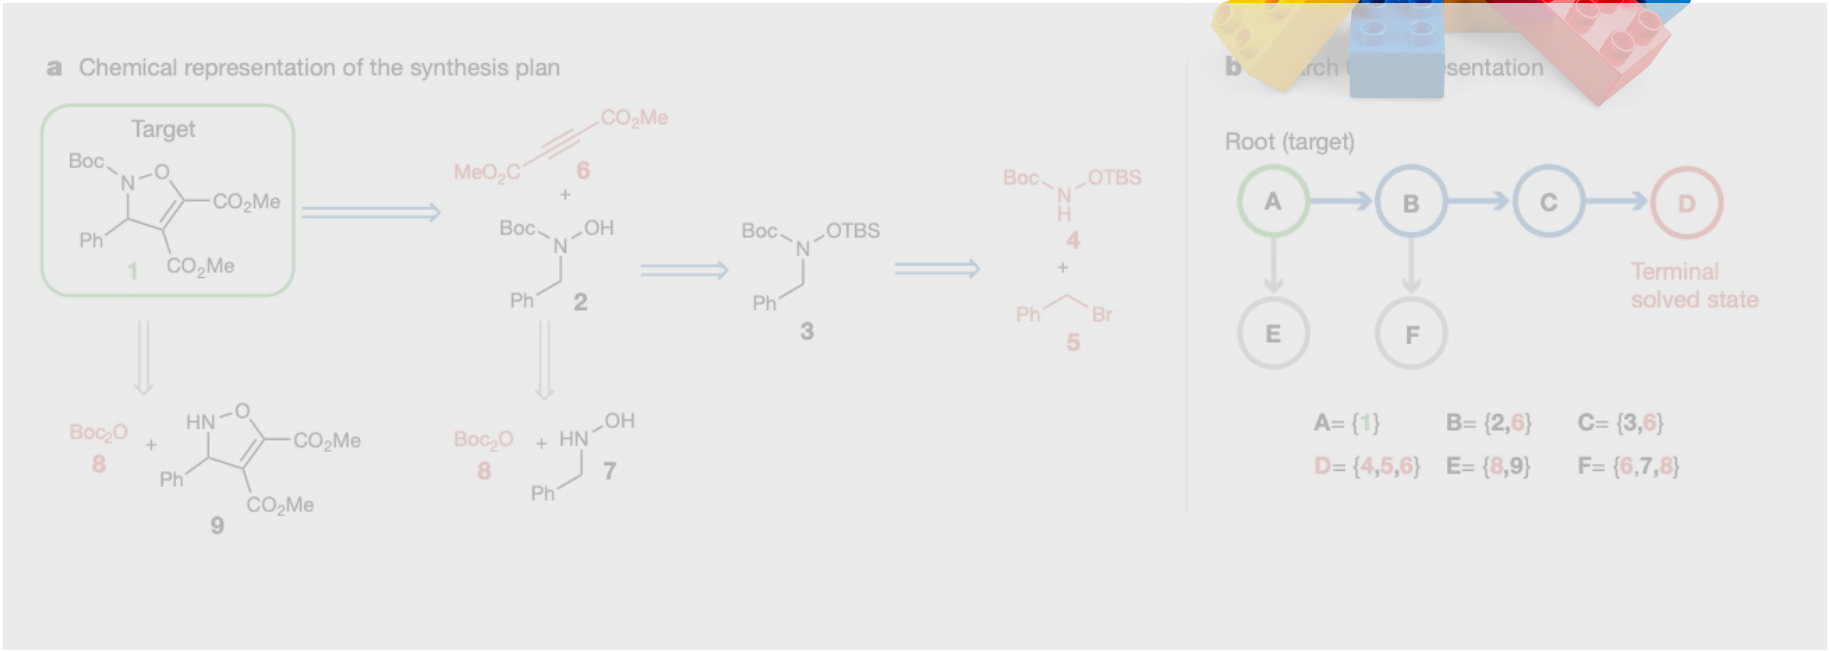

Example retrosynthesis

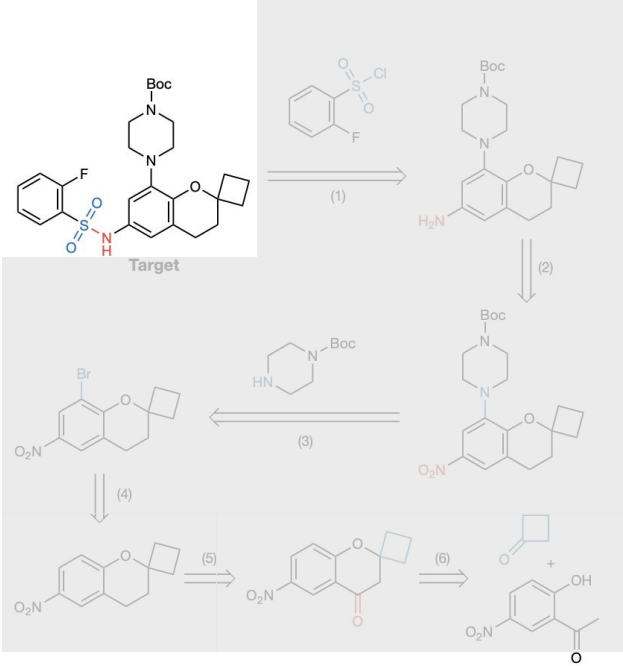

Figure 3 | An exemplary six-step synthesis route for an intermediate in a drug candidate synthesis. This route is identical to the published one[44] and was found by our algorithm autonomously within 5.4 s. The affected functional groups in each step are marked blue or red.

Why is it hard for computers?

- Enormous search space (~10^{60})
- Hand-coded heuristics 60 years ago don’t scale with chemistry literature
- Unlike chess, hard to define “good” intermediate molecules
- Expensive to test if a path works (via wet lab)

Monte Carlo Tree Search

1 Select

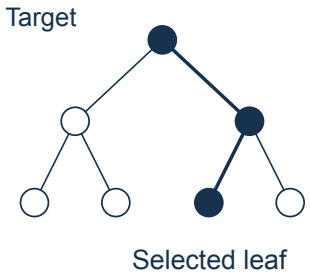

2 Expand

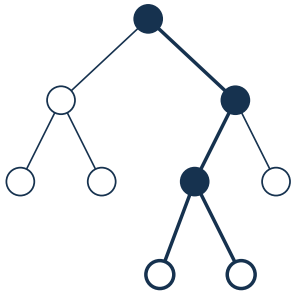

3 Rollout

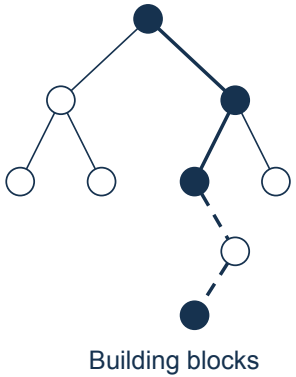

4 Backpropagate

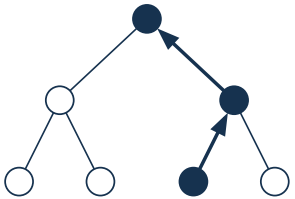

Tree search = memory

Monte Carlo = sampling

MCTS vs UCT vs BFS

MCTS

Search framework

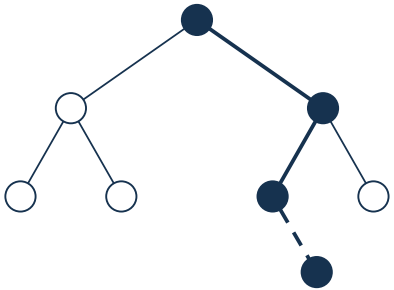

Build a tree and use sampled continuations to evaluate branches.

Upper Confidence Bound (UCT)

A selection rule for MCTS

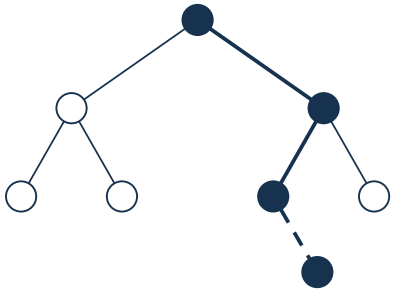

Version of MCTS, branch choice via Exploitation and Exploration

Best-first search (BFS)

Priority-queue search

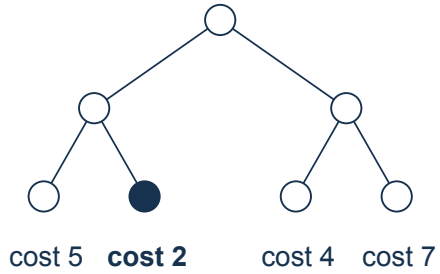

Expand the branch with the lowest heuristic or neural cost.

No rollout evaluation in this paper.

3N-MCTS

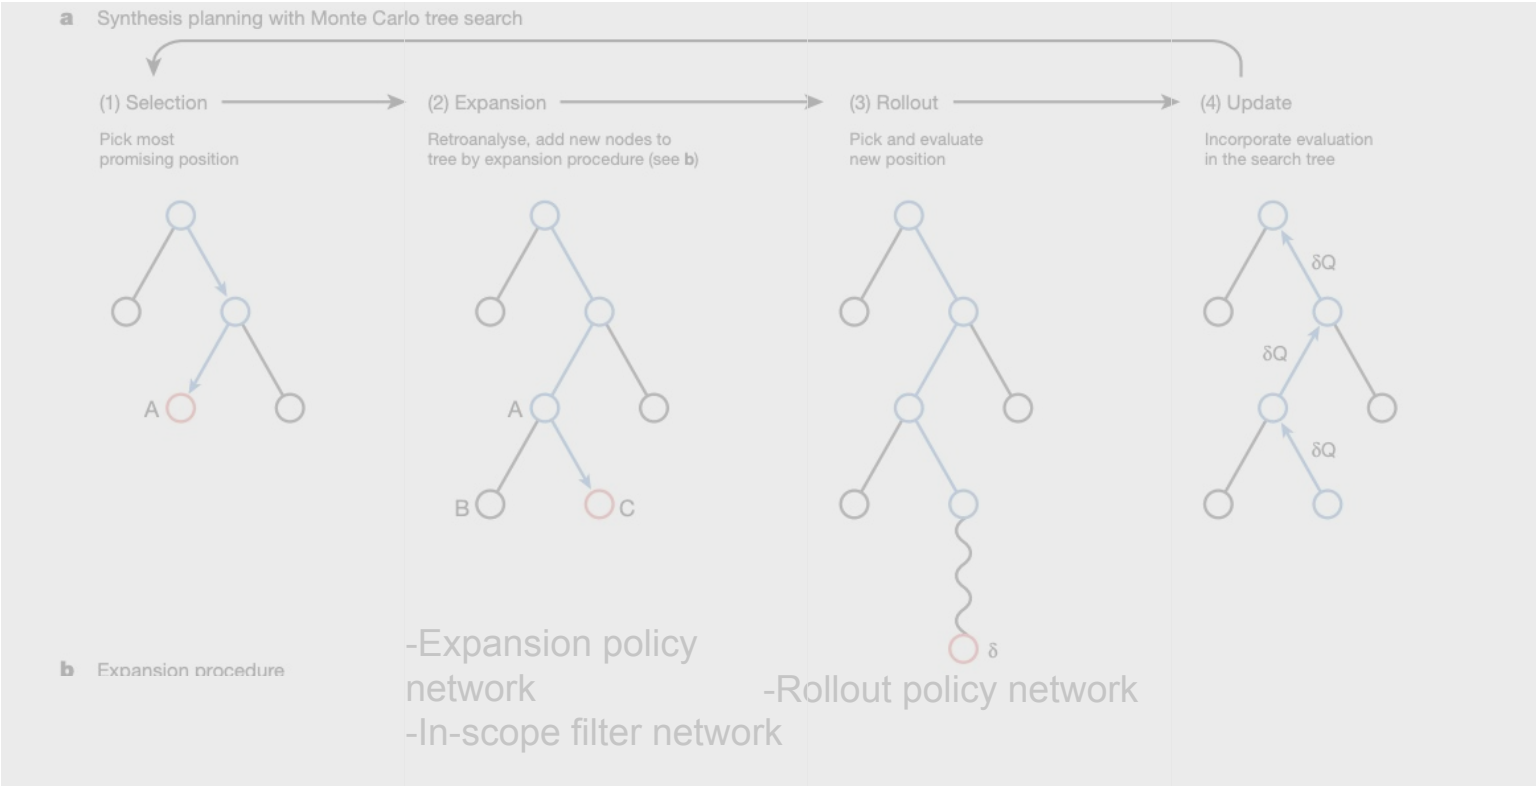
a Synthesis planning with Monte Carlo tree search
(1) Selection
Pick most promising position
(2) Expansion
Retroanalyse, add new nodes to tree by expansion procedure (see b)
(3) Rollout
Pick and evaluate new position
(4) Update
Incorporate evaluation in the search tree
A
A
B
C
δQ
δQ
δQ
δ
b Expansion procedure
-Expansion policy network
-In-scope filter network
-Rollout policy network

Data

- 12.4 million single-step reactions from Reaxys database
- Time-split validation: pre-2015 reactions to train
- Selected # rules from rules occurring >3 (or >50) times

Extended Data Table 1 | Metrics for the supervised neural network policies

Policy	# rules	Coverage	Matching rules/mol[b]	Accuracy[a]	top10Acc[a]	top50Acc[a]
Expansion	301,671	0.79	46,175	0.310	0.633	0.725
Rollout	17,134	0.52	321	0.501	0.891	0.964

Top10Acc/top50Acc is the ratio of correct/incorrect predictions if we allow the system to make 10 or 50 predictions.
[a]Accuracy is calculated on the molecules covered by the respective rulebase.
[b]Matching rules/mol corresponds to the branching factor.

3N-MCTS

$$a_t = \underset{a \in \mathcal{A}(s_t)}{\operatorname{argmax}} \left(\frac{Q(s_t, a)}{N(s_t, a)} + cP(s_t, a) \frac{\sqrt{N(s_{t-1}, a_{t-1})}}{1 + N(s_t, a)} \right)$$

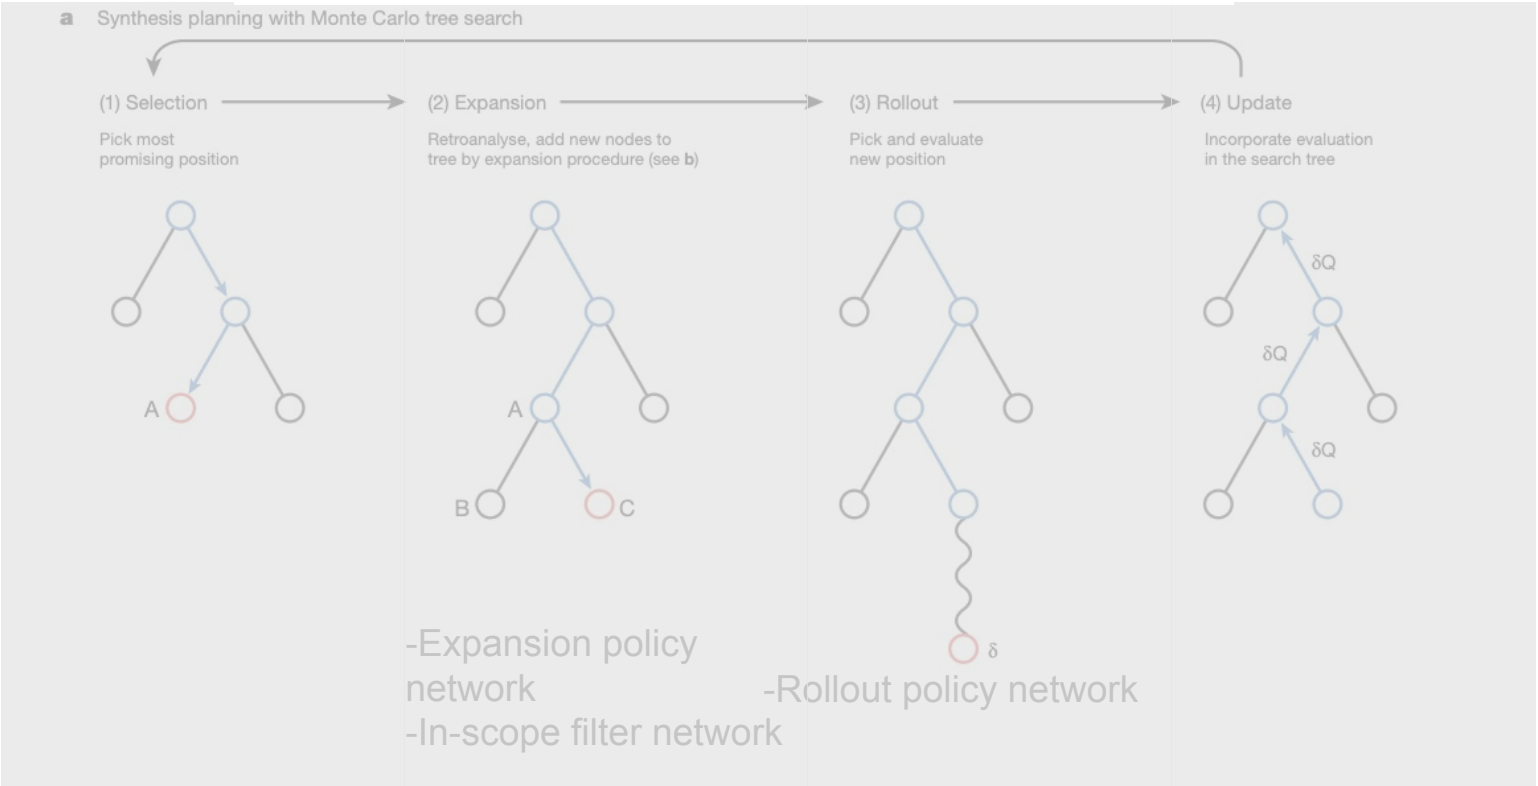

-Expansion policy network
-In-scope filter network

-Rollout policy network

Expansion step mechanics

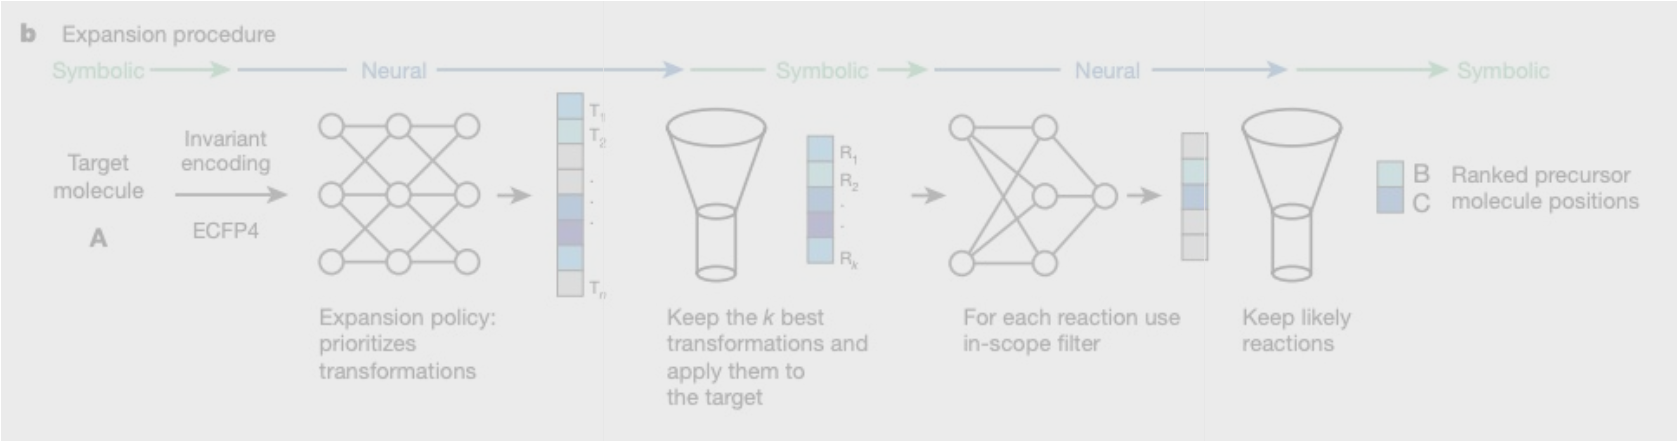

3N-MCTS

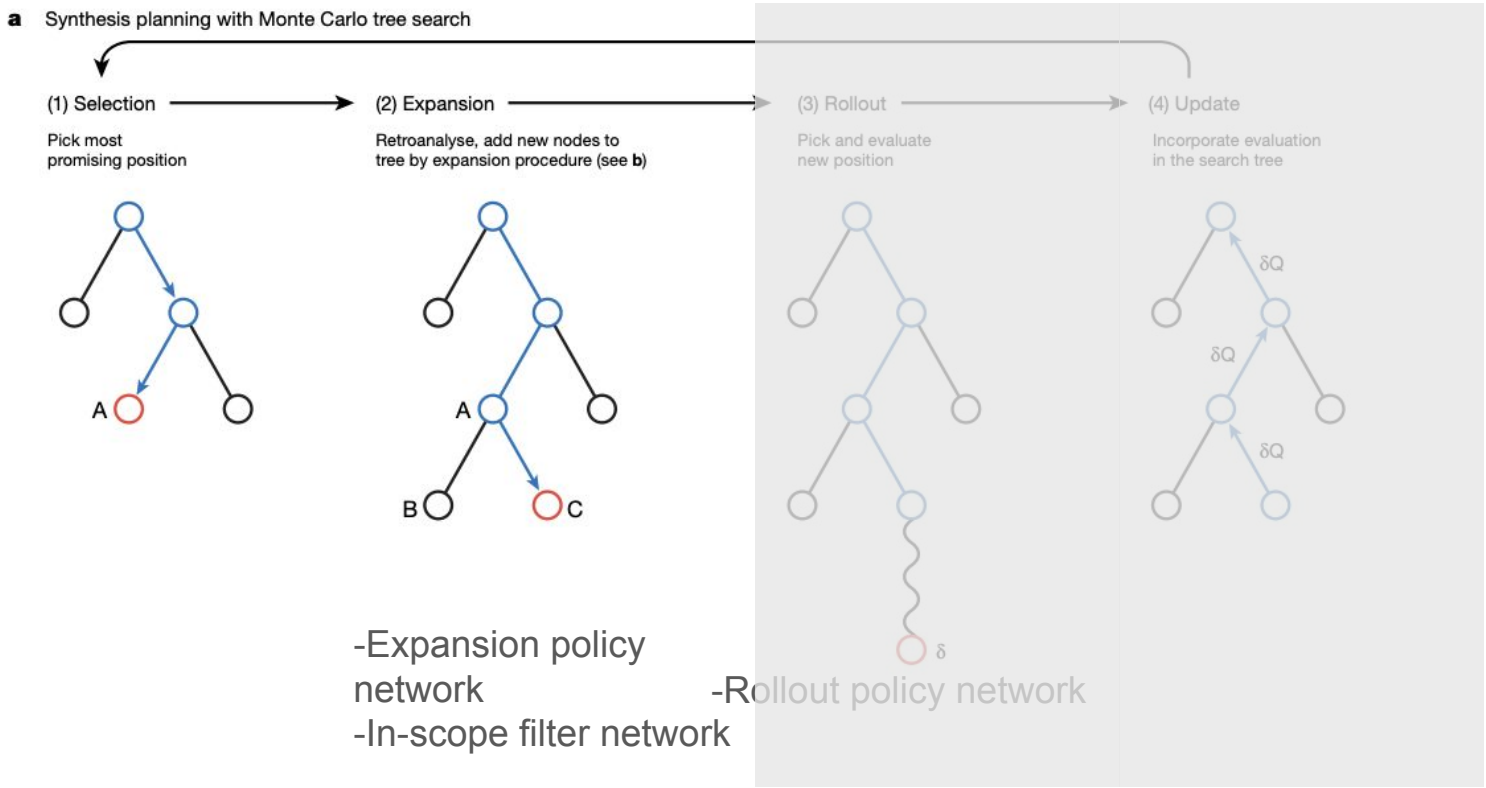

-Expansion policy network
-In-scope filter network
-Rollout policy network

3N-MCTS - Model Breakdown of “3N”

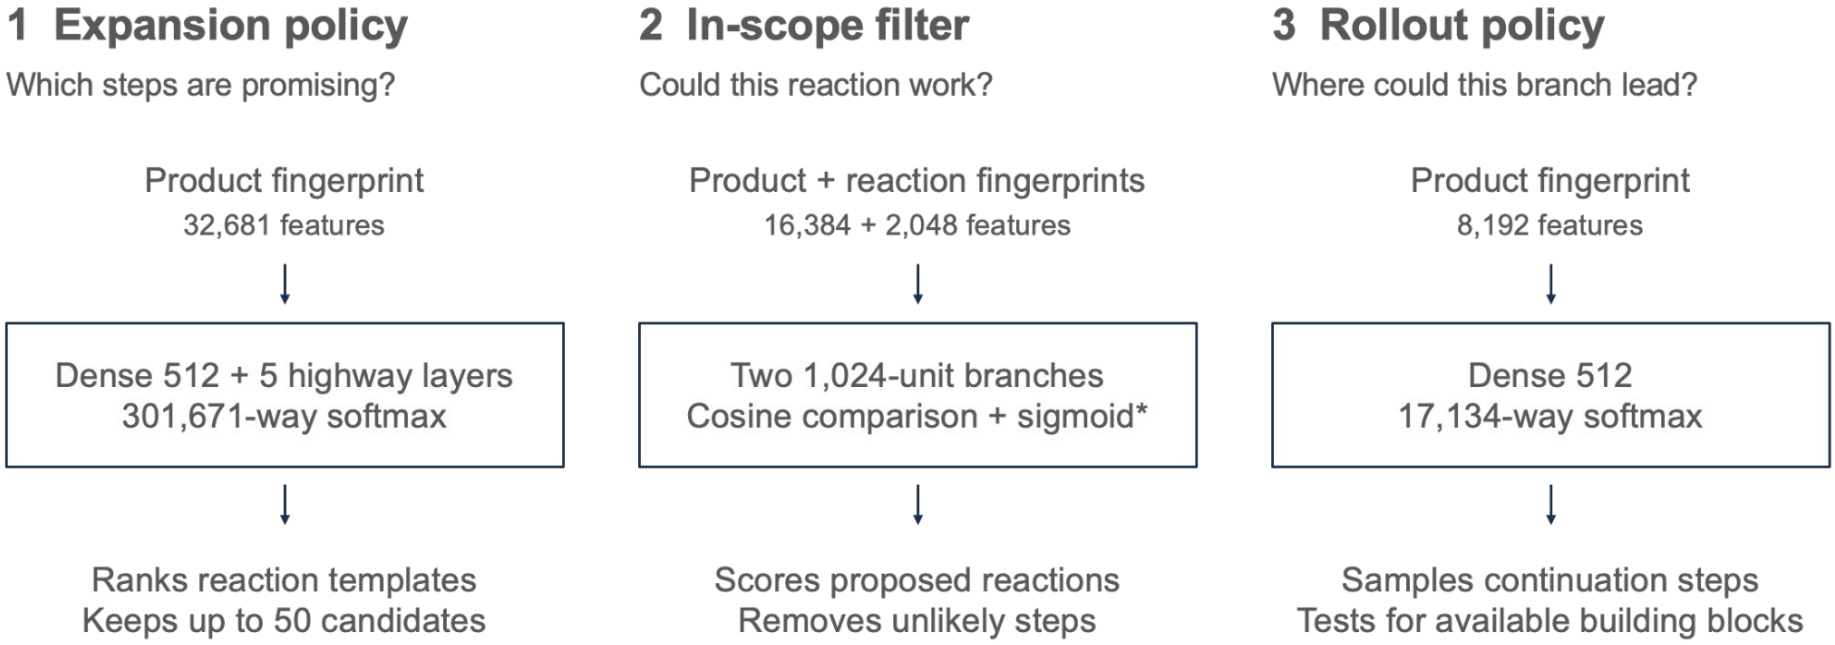

Performance metrics

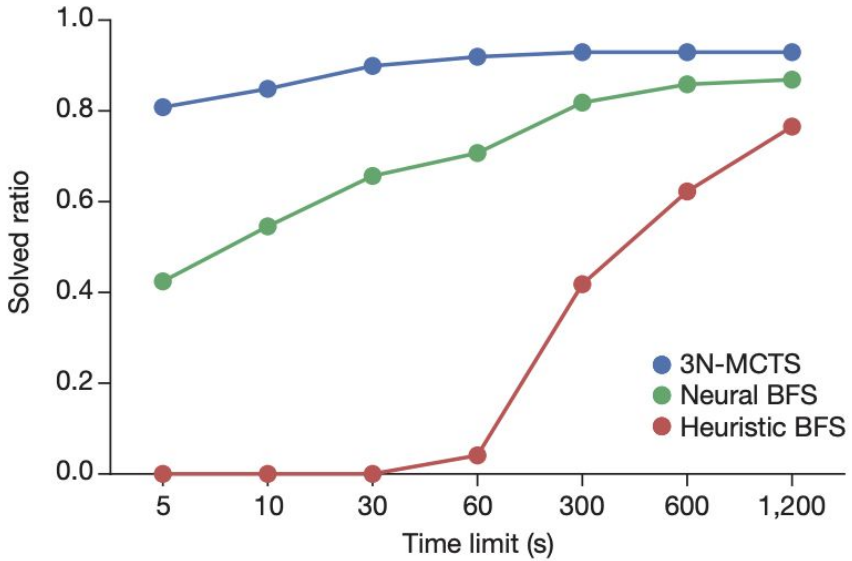

Table 1 | Experimental results

Entry	Search method	Policy*	Percentage solved	Time (seconds per molecule)
1	MCTS	Neural	95.24±0.09	13.0
2	MC	Neural	89.54±0.59	275.7
3	UCT	Neural	87.12±0.29	30.4
4	BFS	Neural	84.24±0.09	39.1
5	BFS	SMILES$^{3/2}$	55.53±2.02	422.1

The time budget was 300 s and 100,000 iterations for MCTS or 300 s and 100,000 expansions for BFS, per molecule. Three restarts were carried out.
*In the BFS, this is the cost function.

- 30x speedup
- 2x targets solved

Double-blind AB tests

- 45 grad level chemists
- Assigned 9 or 10 tasks
- Neither participants nor conductors knew which route came from literature, MCTS, or BFS

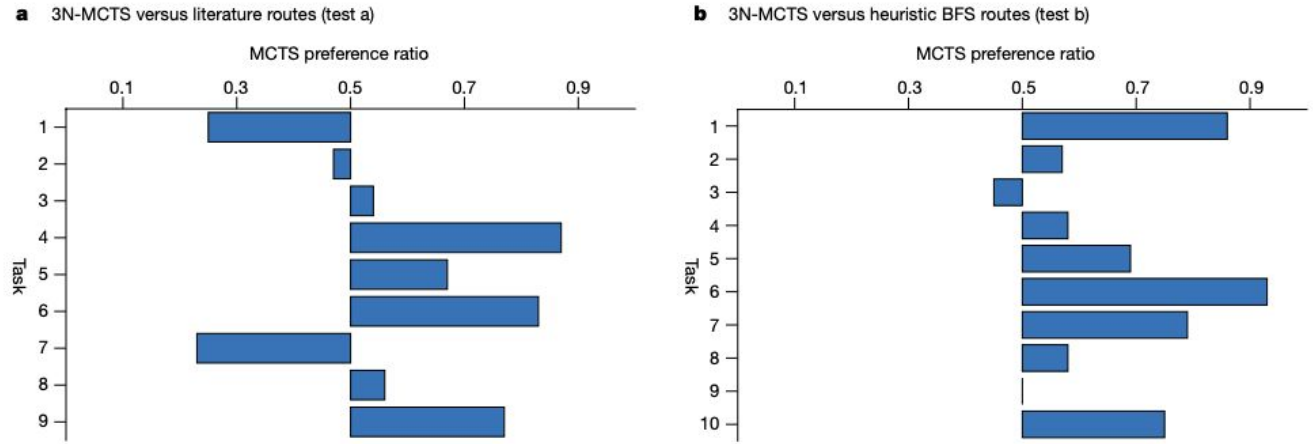

In-scope filter (environment)

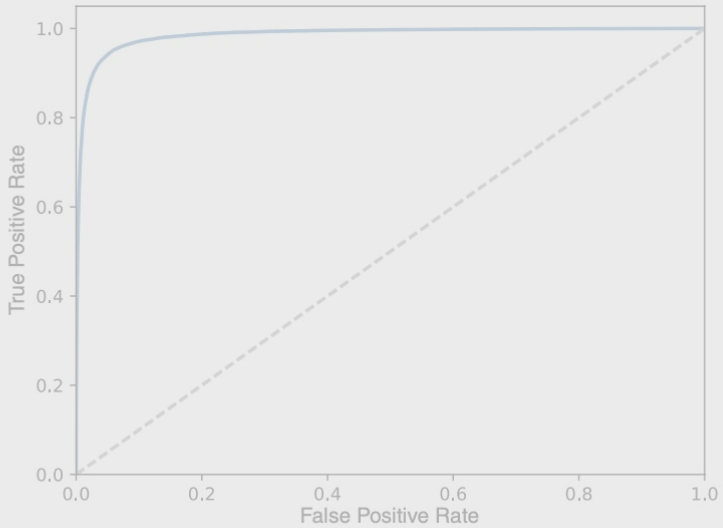

Extended Data Figure 1 | Receiver operation characteristic curve for the in-scope filter. The area under the curve is 0.99.

varying the threshold

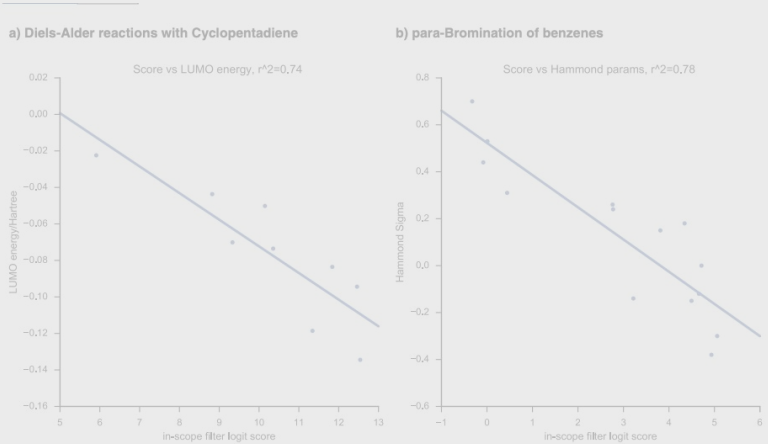

Extended Data Figure 5 | Rediscovering physicochemical properties with the in-scope filter. The output logit score of the neural network correlates surprisingly well with calculated quantum-mechanical properties (LUMO energies, in Hartree) in Diels–Alder reactions ($r^2 = 0.74$) (**a**) and with empirically measured Hammond parameters in electrophilic brominations ($r^2 = 0.78$) (**b**), even though the input features (ECFP4 fingerprints) do not contain electronic information.

learns real physical-chemistry

Limitations

a) Why did chemists prefer the literature over MCTS in Test a) Task 7?

MCTS-Generated key step

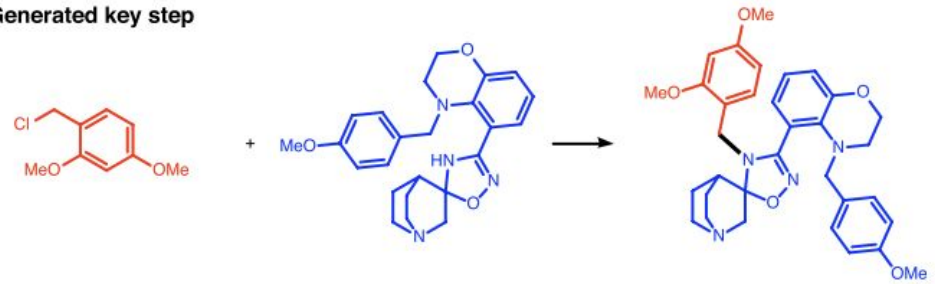

Literature key step

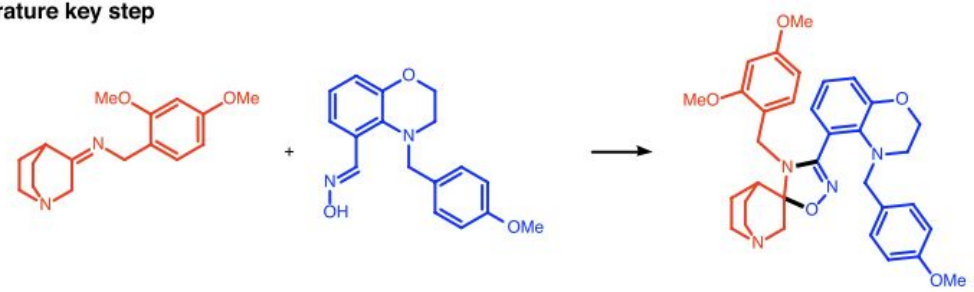

Relevant works prefacing Segler et al.



- learned policies
- fast rollouts
- Use of MCTS



- Neural networks can rank transformation rules and molecular-context effects

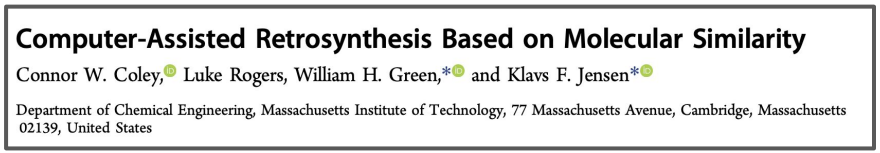

- Demonstrates similarity-based approach for retrosynthesis (RetroSim)

Work that happened after this paper

Article | Published: 21 September 2026

Chemist-aligned retrosynthesis by ensembling diverse inductive bias models

Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy†

Philippe Schwaller, *[a] Riccardo Petraglia,[a] Valerio Zullo,[b] Vishnu H. Nair,[a] Rico Andreas Haeuselmann,[a] Riccardo Pisoni,[a] Costas Bekas,[a] Anna Iuliano [b] and Teodoro Laino[a]

Article | Open access | Published: 25 May 2023

Retrosynthesis prediction using an end-to-end graph generative architecture for molecular graph editing

Weihe Zhong, Ziduo Yang & Calvin Yu-Chian Chen

What did this paper teach us?

- Neural nets and MCTS + policy works
- Neural nets can learn physical-chemical properties
- Newer technology works

Discussion prompts

- Are/will LMs, diffusion models be good at retrosynthesis?
- How far can RL for chemistry go when true environment requires wet lab (robotics?)