doi:10.1038/nature25978

Supporting Information for Planning Chemical Syntheses with Deep Neural Networks and Symbolic AI

M.H.S. Segler, M. Preuss, M. P. Waller

marwin.segler@uni-muenster.de, waller@shu.edu.cn

Contents

1 MCTS Parameter Study by Design of Experiments

1.1 Methods

To assess the robustness of MCTS with respect to the choice of parameters, design of experiments (DOE) was employed. Here, we focused on four key parameters: Time limit per search, maximal depth of the rollout, exploration constant, and the reward given when a solution is proven. A maximin reconstruction (a sophisticated Latin Hypercube Sampling (LHS) [1] variant defined in [2]) with 200 data points and no repetition was performed.

Figures 1 and 2 show the results of that study in a parallel coordinate plot, with parameter ranges as listed in Table 1. As the test set, a randomly sampled set of 500 drug-like molecules first synthesized in 2015 or later was used.

Table 1: Studied MCTS Parameters

Parameter	min	max	scale
Exploration c	0.085	9.8	linear
Rollout depth	0	10	linear
Reward when proven	1	99.52	log
Search time/s	5	1200	log

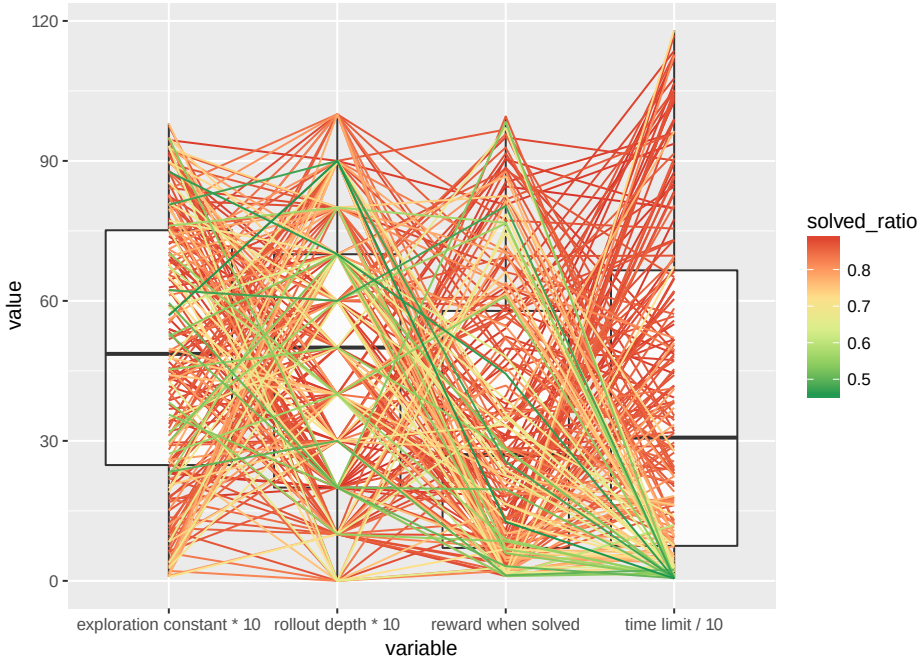

Figure 1: Results of the full DOE (200 points). Parameters are scaled component-wise as indicated, the best configurations are marked in red. For the first two parameters, we have used a linear scaling, for the two latter, a logarithmic scaling, meaning fewer points with high values.

1.2 Results and Discussion

In the study, MCTS turned out to be quite robust over a wide range of values for the the exploration constant, the proven reward and the rollout depth. The most important parameter is search time. Not surprisingly, shorter search times lead to a lower solution ratio.

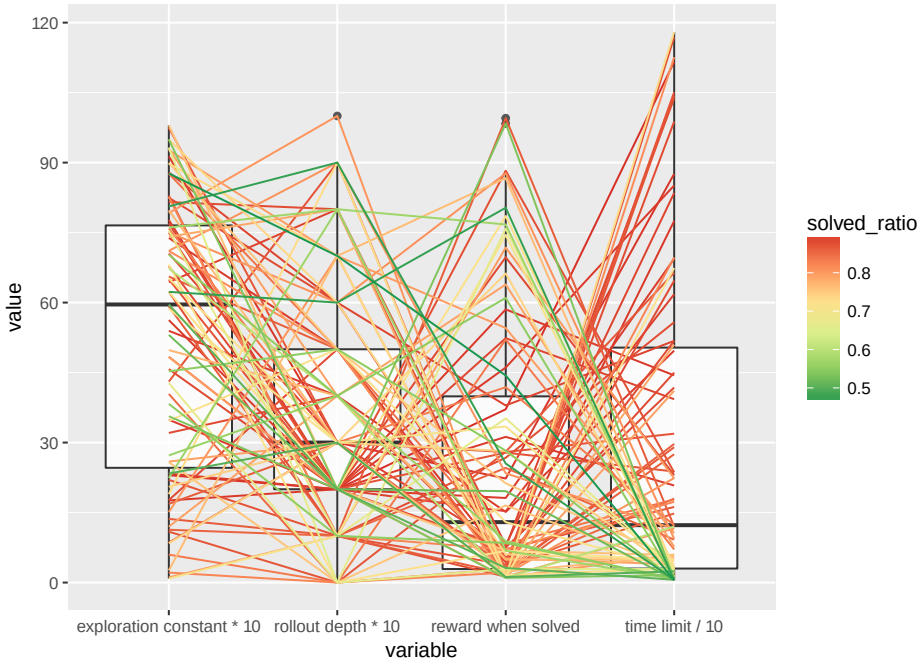

Figure 2: Results of the DOE, only Pareto optimal points regarding the three objectives solved ratio, mean depth, mean time are shown. Parameters are scaled component-wise as indicated, the best configurations are marked in red.

2 Analysing Route Diversity

To assess route diversity, we retrieve 19 additional routes besides the route with highest scored. For 3N-MCTS, 600 s total simulation time was allocated for all 20 routes, while neural BFS received a budget of 600 s per retrieved route. Even though this means that BFS received 20 times more search time than MCTS, it retrieved less routes than MCTS, often not retrieving any routes at all.

2.1 Depiction Method

To depict the routes, counted AtomPair fingerprints with a maximum distance of four, folded to a length of 1024, and $\log(x+1)$-preprocessed, were calculated for each state using rdkit. These features were subjected to Principal Component Analysis (PCA), and the first two principal components were used to depict the states in 2D. In the figures below, the routes found by BFS are coloured red, and MCTS blue. The solutions (terminal solved states = the set of building blocks) are denoted with two numbers n-m, where n is is the target molecule and m is the index number of the route, start with 0 as the best route. E.g. 3-0 is the best route for the third target molecule, 5-2 is the third-best route for the fifth target molecule.

2.2 Results and Discussion

Both methods tend to find solutions in similar regions of chemical space, with the difference of MCTS being two orders of magnitude faster. In some cases, one of the methods finds more diverse routes than the other, but there is no consistent trend across the different targets. Interestingly, both methods tend to explore different solutions only later in the tree.

To encourage the exploration of diverse routes, one could for example submit not just the root node to the search, but also the first layer of children of the root individually. Alternatively, scoring functions could be introduced to discourage the exploration of routes similar to those already returned, or routes could be ordered by cluster membership, after clustering the solution states.

2.3 Example Routes

Route 1

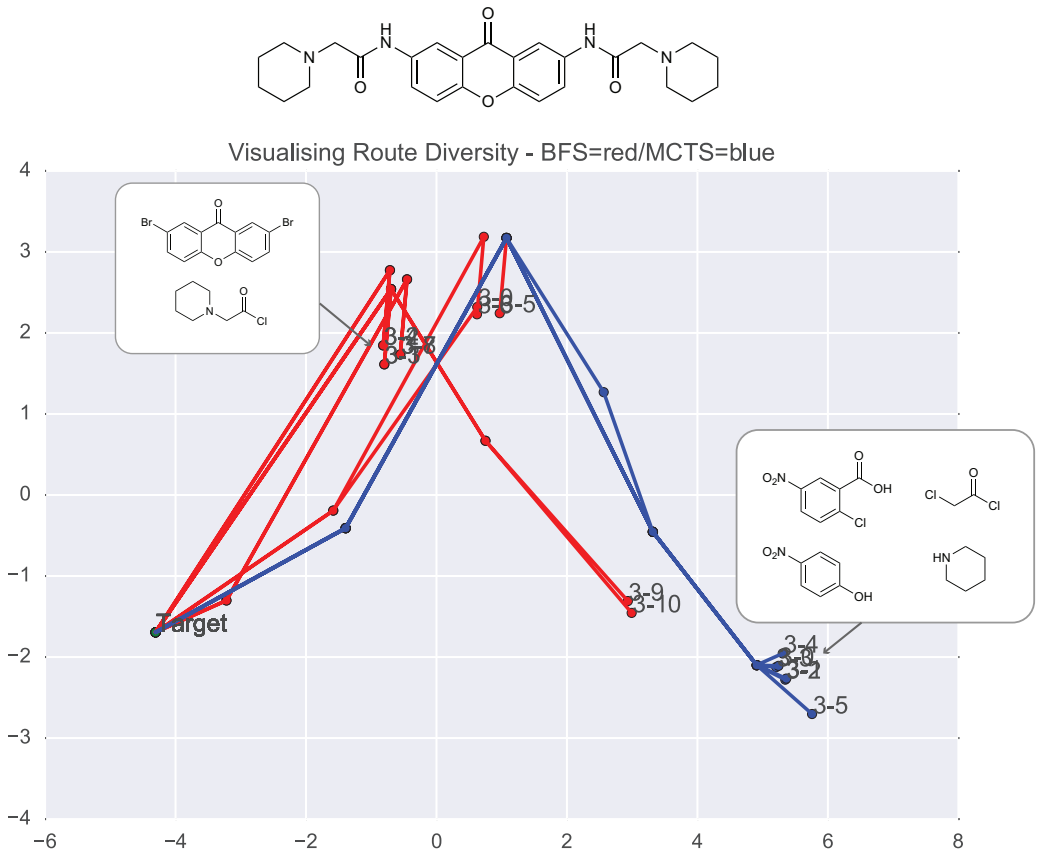

Figure 3: Route 1

Route 2

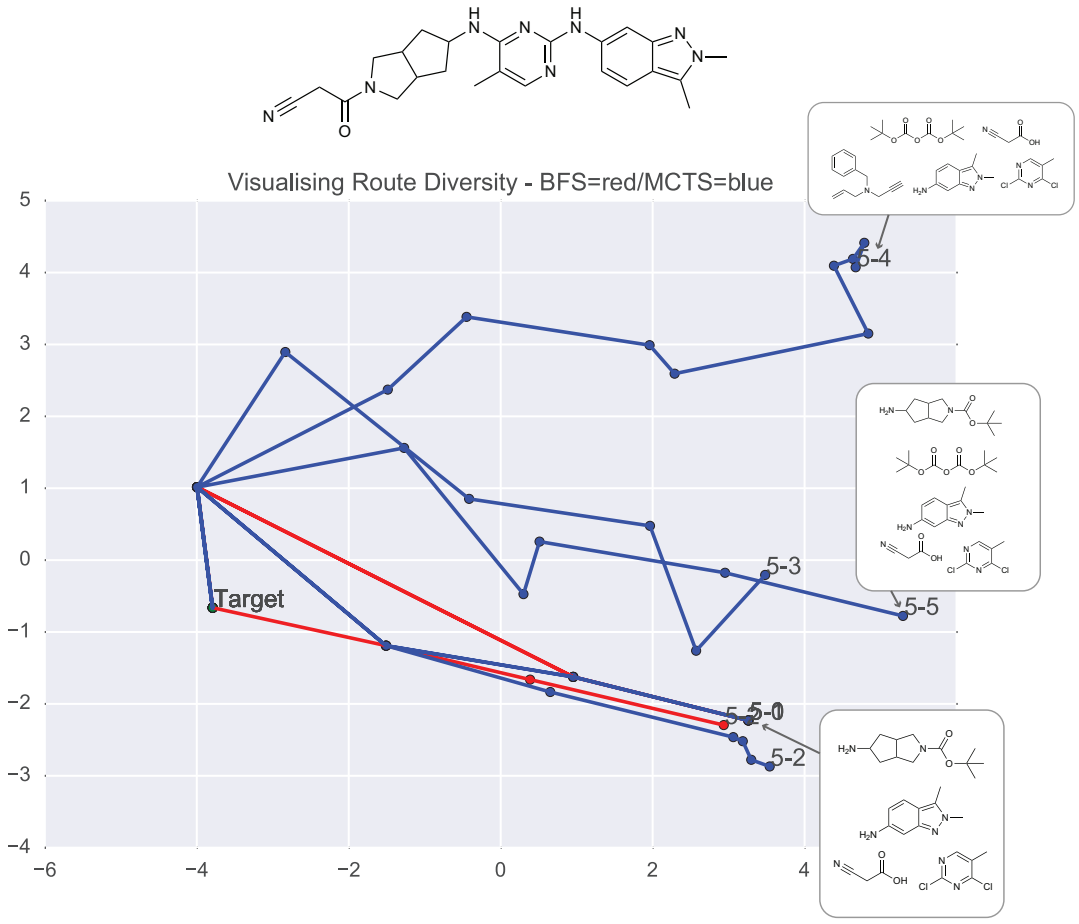

Figure 4: Route 2

Route 3

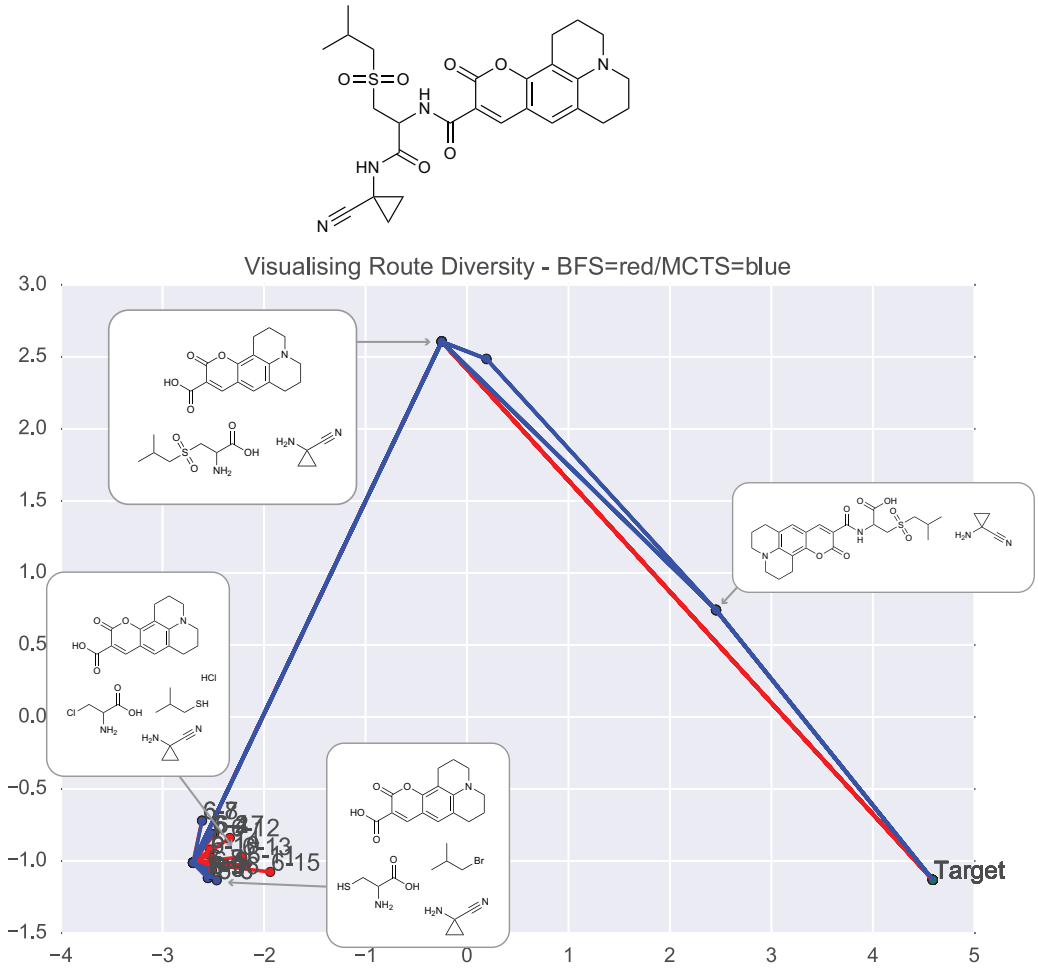

Figure 5: Route 3

Route 4

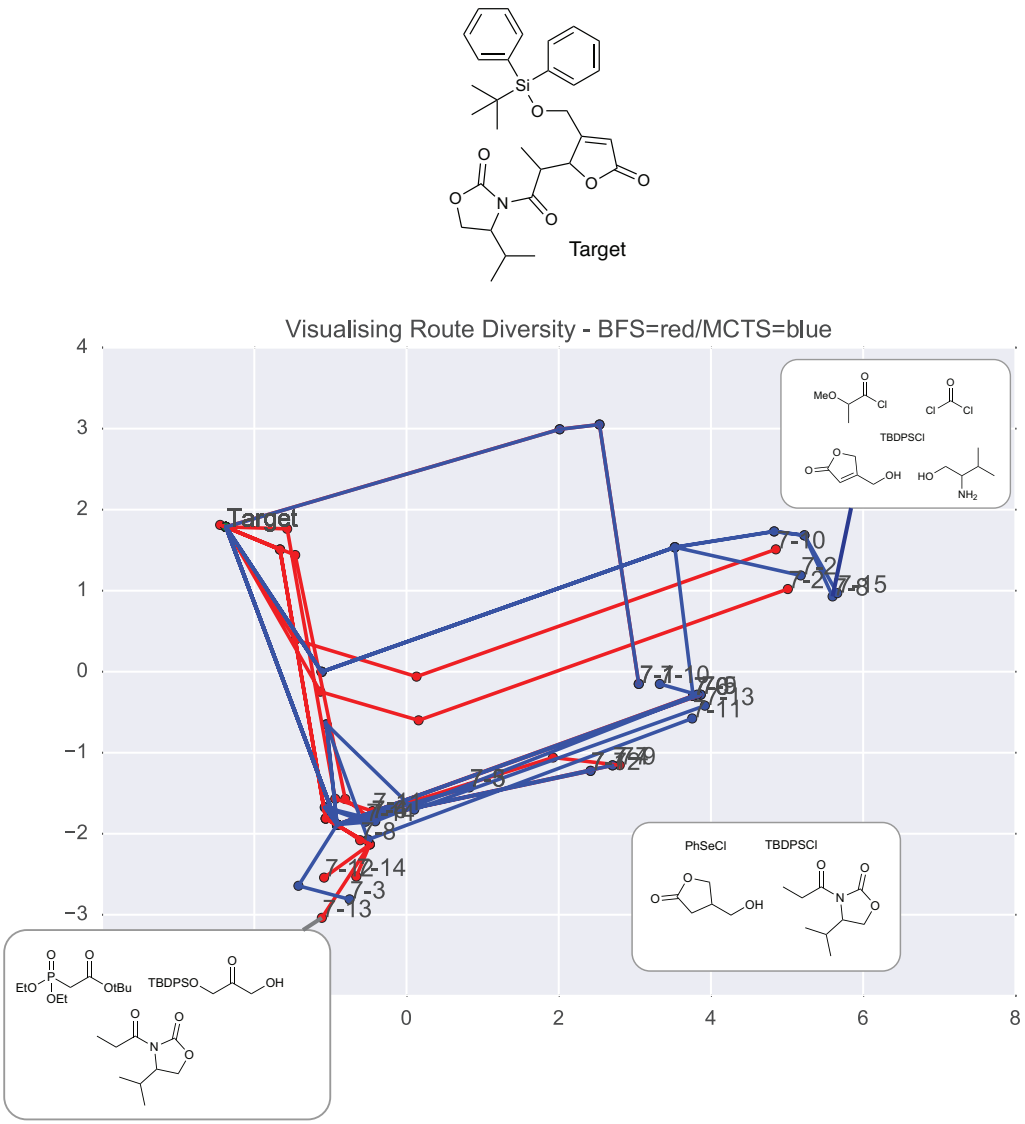

Figure 6: Route 4

Route 5

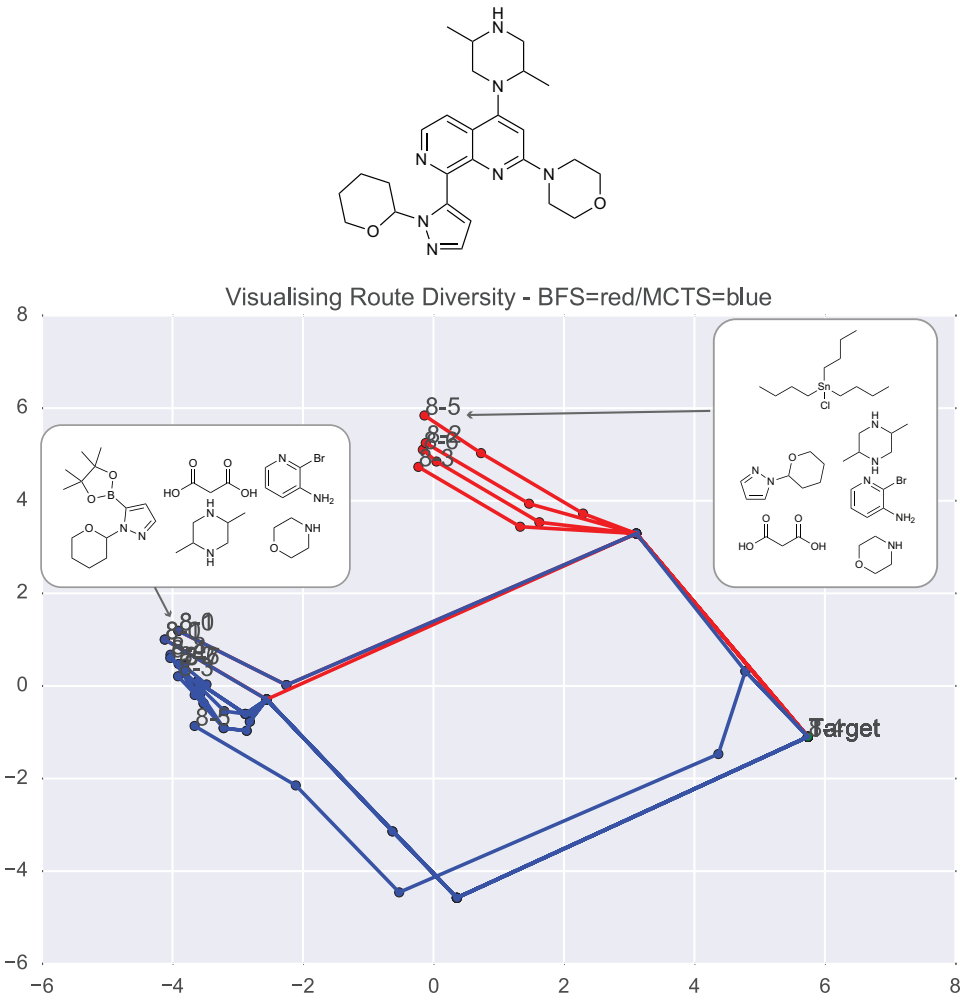

Figure 7: Route 5

3 Unsolved Target Molecules from our Quantitative Study

Below are 7 randomly selected molecules that could not be solved with either the 3N-MCTS algorithm nor the other algorithms (BFS or neuralBFS) from the targets used in the quantitative study.

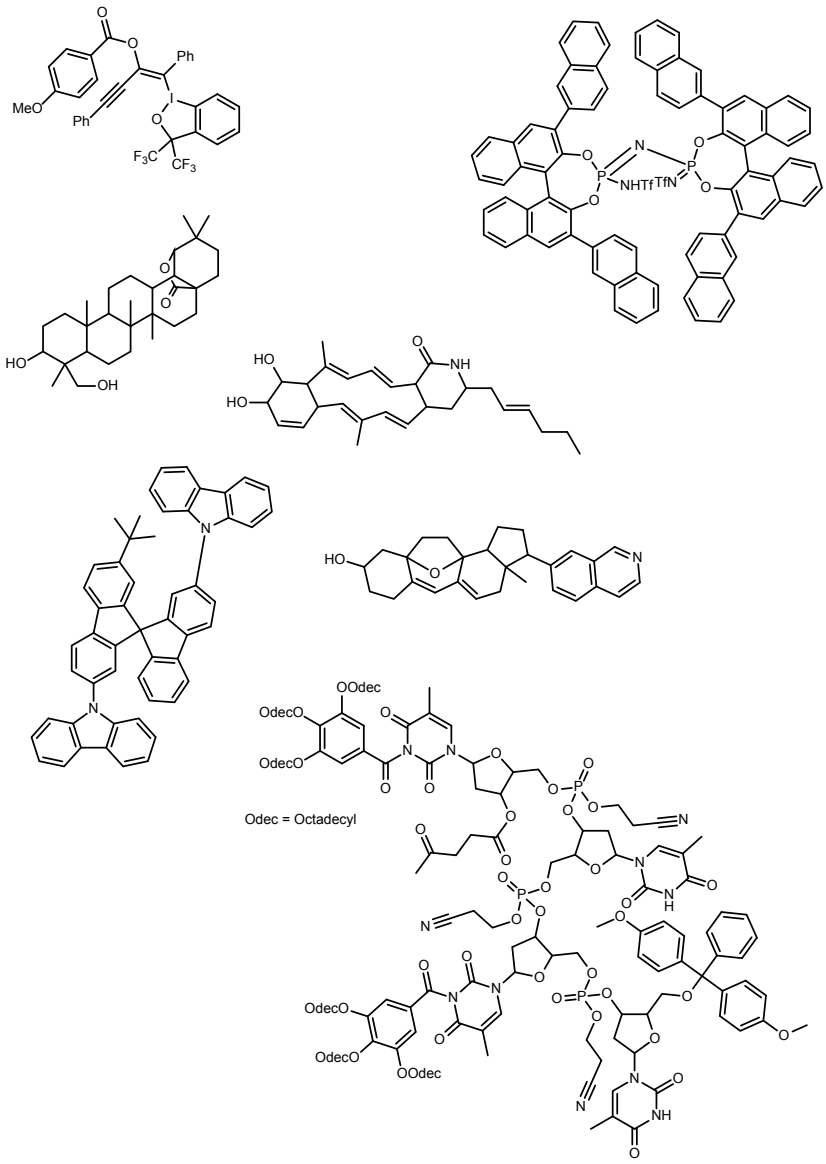

References

[1] McKay, M. D., Beckman, R. J. & Conover, W. J. Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**, 239–245 (1979).

[2] Wessing, S. *Two-stage methods for multimodal optimization*. Ph.D. thesis, Technische Universität Dortmund (2015). URL `http://hdl.handle.net/2003/34148`.