Appendix

A Some Additional Details on the Dissimilarity Measure

A.1 Solving (1)

In this section, we describe how the linear program for computing the dissimilarity measure (1), can be solved using an optimal transport (OT) program Villani (2008). This reformulation is similar to that of Kandasamy et al. (2018), who use OT to describe a distance between neural network architectures.

Say we are given two molecules $M_1 = (A_1, B_1), M_2 = (A_2, B_2)$ with n_1, n_2 atoms respectively, let $U \in \mathbb{R}_+^{n_1 \times n_2}$ denote the matching matrix, i.e. $U(i,j)$ is the weight matched between $i \in M_1$ and $j \in M_2$. We now define a sequence of variables which form the parameters of our OT program. First, let $w_{\mathsf{m}}(M_i) = \sum_{a \in A_i} w_{\mathsf{a}}(a)$ is the total weight of a molecule $M_i = (A_i, B_i)$ for $i = 1, 2$. Denote $y_1 = [\{w_{\mathsf{a}}(a)\}_{a \in A_1}, w_{\mathsf{m}}(M_2)] \in \mathbb{R}^{n_1+1}$ and $y_2 = [\{w_{\mathsf{a}}(a)\}_{a \in A_2}, w_{\mathsf{m}}(M_1)] \in \mathbb{R}^{n_2+1}$. Next, let $C = C_{\text{at}} + C_{\text{st}} \in \mathbb{R}^{n_1 \times n_2}$ and $C' = [C\, \mathbf{1}_{n_1}; \mathbf{1}_{n_2}^\top\, 0] \in \mathbb{R}^{(n_1+1)\times(n_2+1)}$; i.e. C' has $C_{\text{at}} + C_{\text{st}}$ in its first $n_1 \times n_2$ block, representing the atom type and bond type penalties in (1), while the 1's in the last row and column capture the non-matching penalty. We finally let $U' \in \mathbb{R}^{(n_1+1)\times(n_2+1)}$ be our optimization variable where the first $n_1 \times n_2$ block will correspond to the optimization variable U in the original program. It is easy to see that (1) is equivalent to the following linear program, which is an optimal transport program:

$$\begin{aligned} &\text{minimise} \quad \langle U', C' \rangle \\ &\text{subject to} \quad U' \mathbf{1}_{n_2+1} = y_1, \quad U'^\top \mathbf{1}_{n_1+1} = y_2. \end{aligned}$$

We refer the reader to Theorem 2 in Kandasamy et al. (2018), who formally prove this result in a similar setting.

A.2 T-SNE visualizations for the OT distance

We perform another experiment to verify the validity of the proposed optimal transport dissimilarity measure. We use the four different base combinations of settings for the OT distance to compute distances between 200 randomly sampled molecules, and use these distances to compute 2-dimensional t-SNE embeddings (Maaten and Hinton, 2008). These embeddings aim to preserve distances, so that visual closeness translates into OT-distance closeness. We also color the points by values of QED (drug-likeliness) and synthetic accessibility scores. The results are shown in Figure 10. We see that despite the fact that the chemical space has complicated dependencies between molecule structure and properties, dependencies in the induced embedding space are relatively continuous. We can also observe clusters of molecules with similar values. In Figure 11, we compare the planar embeddings produced by other possible distances: ℓ_2 distance between pairs of fingerprints and inverted Tanimoto similarity measure between molecules (referred to as fingerprint kernel in the main part of the paper), one may say OT-dist looks slightly better (e.g. low versus high values are more separated in the plots).

A.3 Some Known Limitations

Stereoisomers: Since our dissimilarity measure is based on the graph representation, it will not be able to distinguish between stereoisomers, i.e. molecules which have the same formula and bonded atoms, but different 3D orientation. For example, pictured below are D-Glucose and L-Glucose. Since, they have the same graph representation, our dissimilarity measure will be 0 between both molecules. However, they have different 3D structures (being mirror images of each other), which can give rise to different physical properties. For instance, D-Glucose can be digested by the human body while L-Glucose cannot.

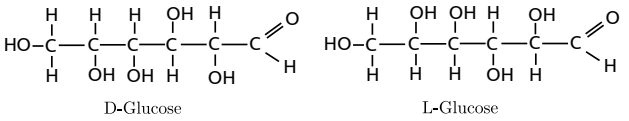

D-Glucose L-Glucose

It is worth noting that many graph convolution based approaches for modeling molecules face this challenge. One way to circumvent this issue is to combine our kernel with other features which account for 3D structure in a sum or product kernel.

B Some Implementation Details

For the BO methods, we fit GP hyperparameters by maximizing the marginal likelihood. As the acquisition, we adopt the ensemble method described in (Kandasamy et al., 2019a) using the EI, UCB, and TTEI acquisitions instead of sticking to a single acquisition. To optimize the acquisition, we ran the explorer for 20 iterations on each BO iteration, but added the new molecules to our initial pool $\mathcal{S}$ for the next iterations, so that we can search across a large pool during the entire optimization routine. This corresponds to "reusing" explored and synthesized compounds in a real experiment.

C Additional Experimental Results

Experiments with low starting value To verify that ChemBO successfully optimizes the objective regardless of the quality of initial pool, we conduct an experiment on pools of 20 molecules randomly selected from subset of ChEMBL dataset that has value of the objective function capped by 0.7 for QED and 3

for penalized LogP function (approximately 60% percentiles in ChEMBL). The results on Figure 4 show that ChemBO performs well in such cases, too, and does so better than baseline with the same regularities as before (the fingerprint kernel performs worse than ot-dist kernel on QED and better on penalized LogP task).

Synthesis Paths

We visualize the synthesis paths for some of the optimal molecules in Figures 6-9. The boxed molecules are from the initial pool of 20 reagents. In this figure, when arrows from two or more parent molecules point to a child molecule, it means that the child molecule was obtained by reacting the parent molecules.

It is worth mentioning some caveats here. First, we see a few cases of complex molecules being combined to produce a simpler molecule – the most striking example being the one in Figure 7 where two complex molecules are combined to produce Methane (CH_4)[5]. It is more likely that simpler molecules will be available as reagents in a realistic setting. This is an artefact of our initial pool, and we believe that such cases can be avoided by carefully selecting an initial pool. Second, note that in all synthesis paths shown, there are molecules with large rings. Large rings are not necessarily stable, and hence such molecules are hard to synthesize. We believe this could be due Rexgen, and, as mentioned in the main text, when such synthesis predictors become more accurate and reliable, so will the efficacy of our proposed framework.

The red boxes in the molecules are because RDkit's 2D layout algorithm overlays two atoms – which is likely to happen with large molecules.

Some statistics on the ChEMBL Dataset: In Figure 5, we plot the distribution of QED and Pen-logP on the ChEMBL dataset. These values help us understand the success of optimization procedures relative to the average over the dataset from which the starting pool was drawn: the histograms show that the optimized values lie in the highest percentiles of the original dataset.

[5]In reality, Methane was probably just meant to be a by-product of a reaction meant to produce some other molecule.

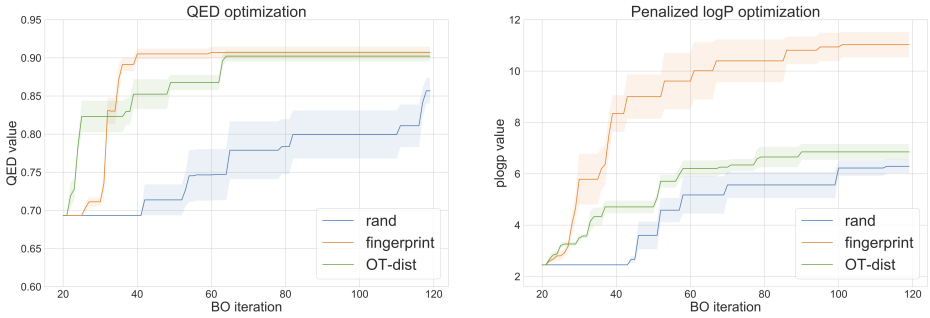

Figure 4: Results comparing the three methods described in the beginning of Section 4. We plot the number of iterations (after initialization) against the highest found QED (left) and Pen-LogP (right) values by each method. Higher is better in both cases. All curves were produced by averaging over 5 independent runs. The shaded regions indicate one standard error.

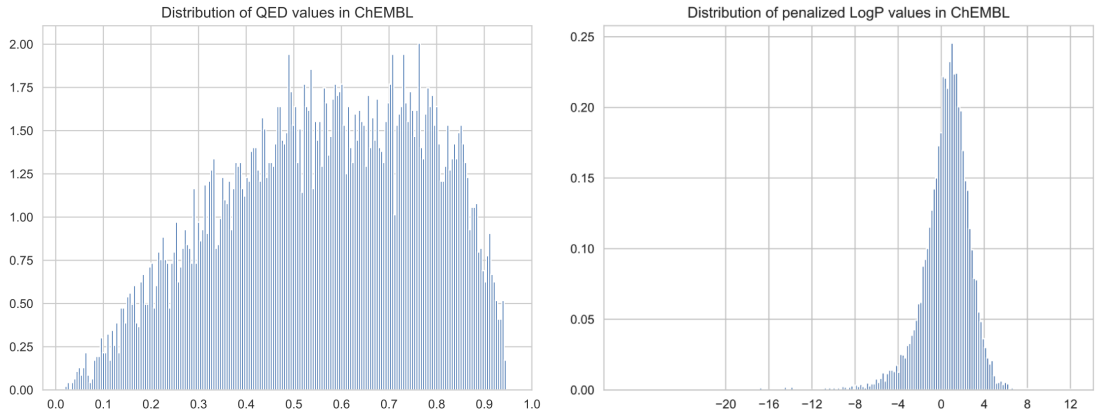

Figure 5: ChEMBL dataset statistics: normalized histograms of QED score and penalized logP score.

Figure 6: Synthesis path for molecule with penalized logP 11.988. The boxed molecules are from the initial pool of 20 reagents.

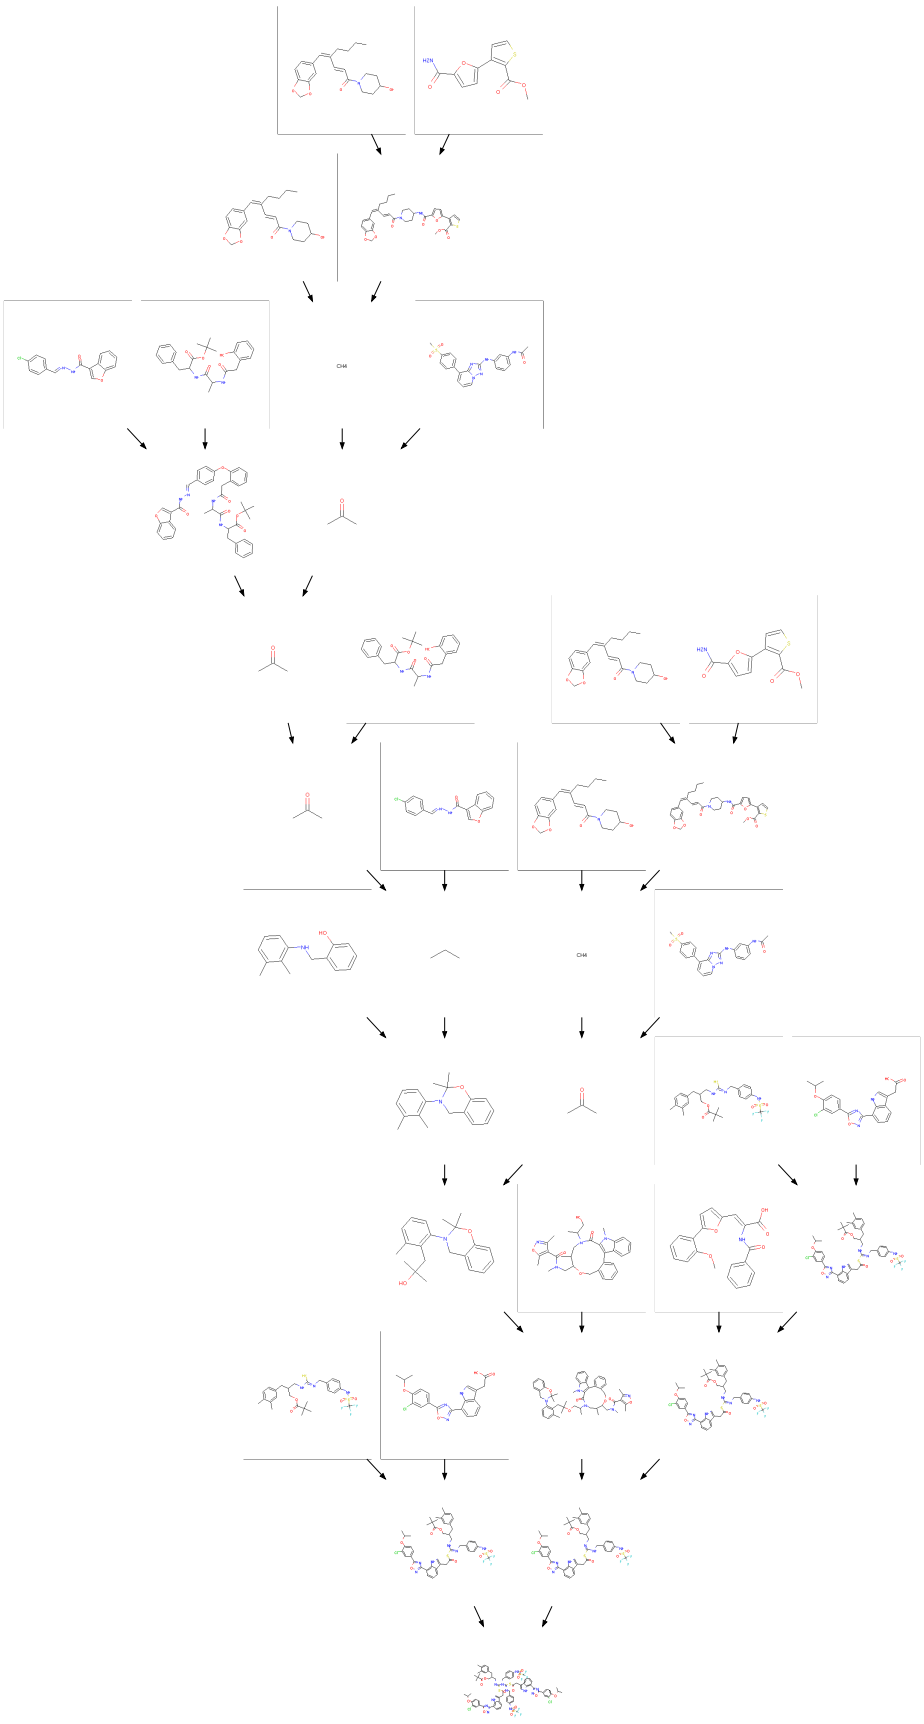

Figure 7: Synthesis path for molecule with QED 0.92. The boxed molecules are from the initial pool of 20 reagents.

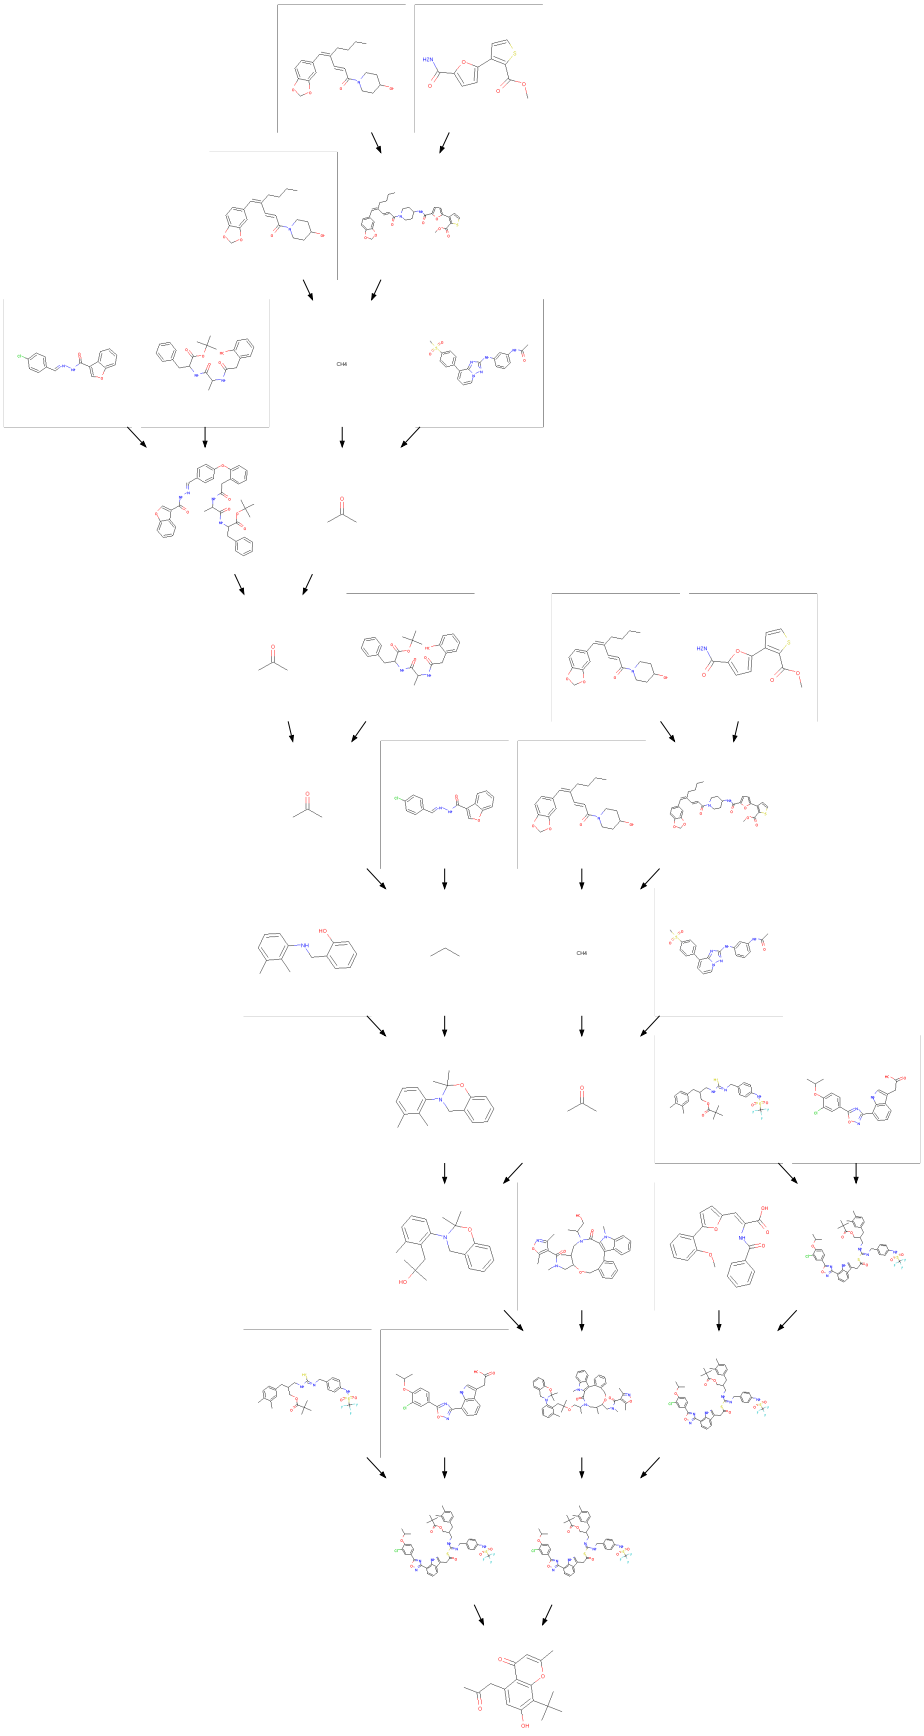

Figure 8: Synthesis path for molecule with penalized logP 8.306. The boxed molecules are from the initial pool of 20 reagents.

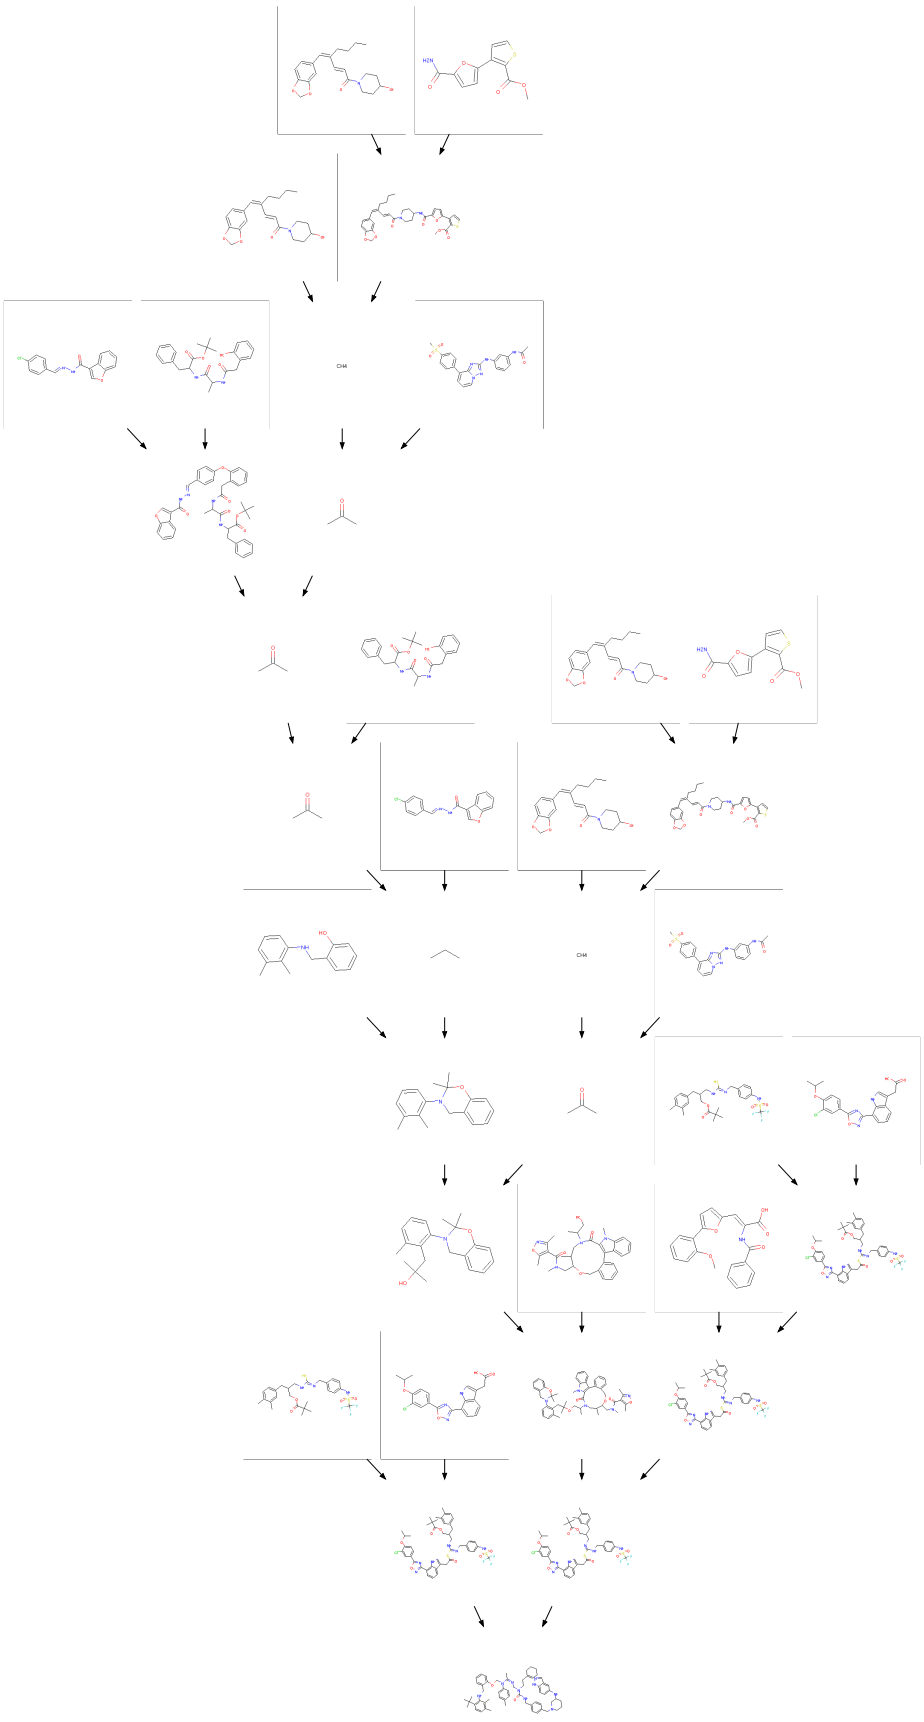

Figure 9: Synthesis path for molecule with QED 0.93. The boxed molecules are from the initial pool of 20 reagents.

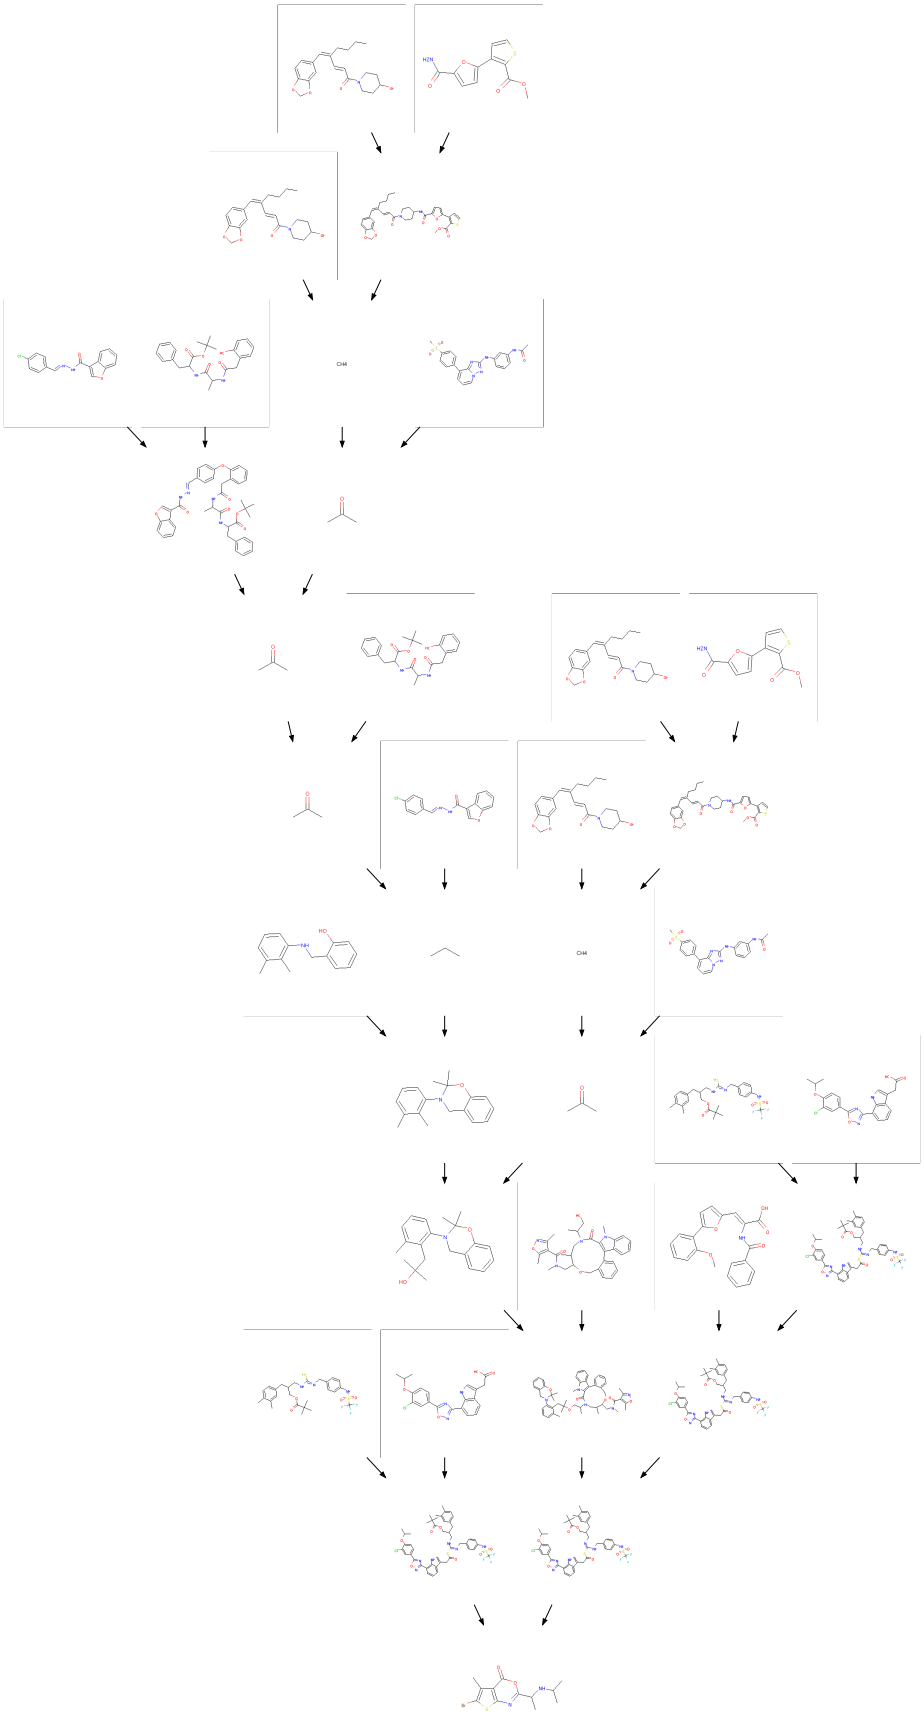

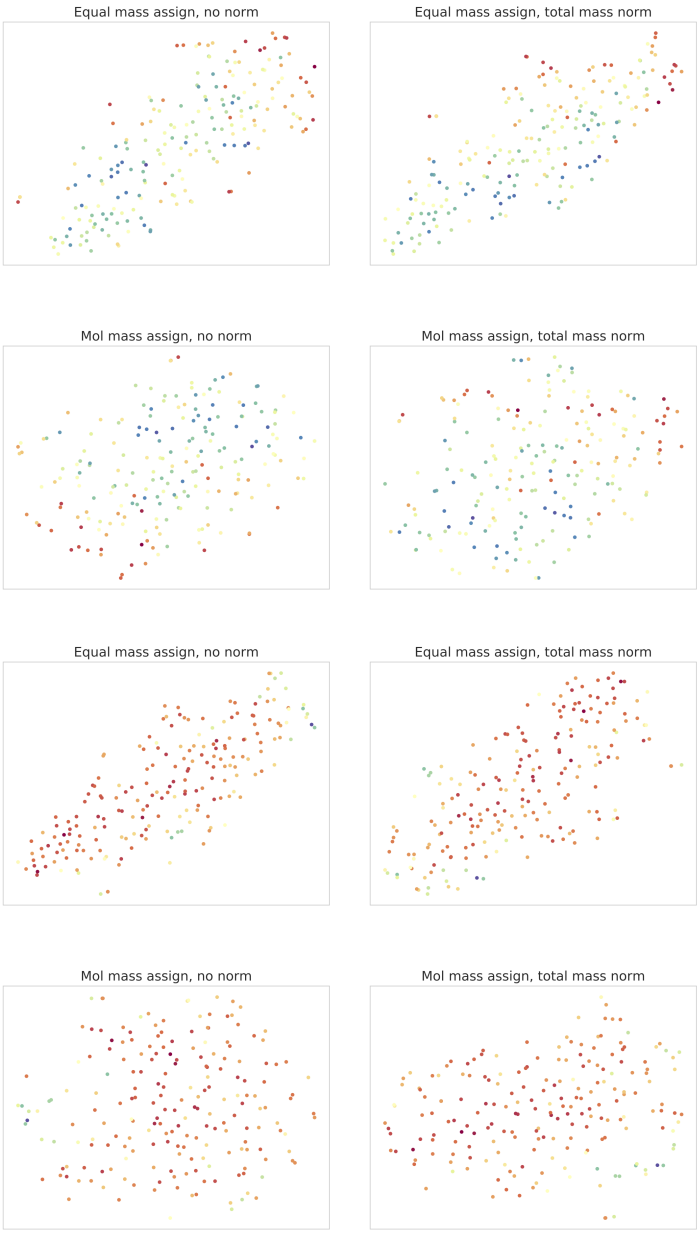

Figure 10: t-SNE visualization of OT distance ot-dist for different parameter configurations, first four color-coded by QED value, last four by SA score.

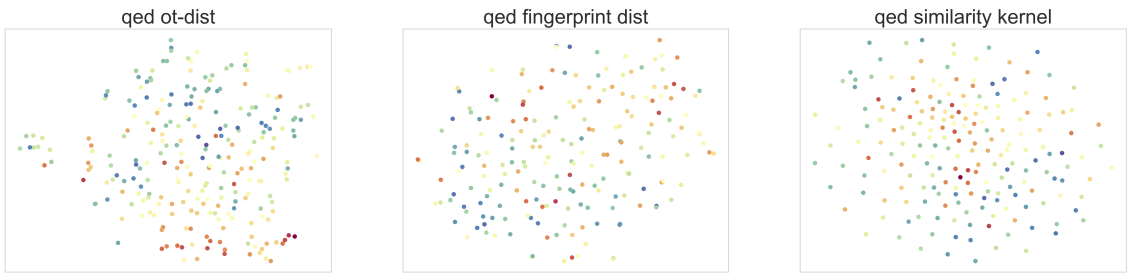

Figure 11: Comparison of t-SNE embeddings produced based on three molecular distances: ot-dist, ℓ_2 distance between fingerprint vectors, and inverted similarity kernel between fingerprints.