

Scaffold-tuner: a property-directed compound editing tool for scaffold-preserving structural optimization

Tsuyoshi Esaki

`tsuyoshi-esaki@biwako.shiga-u.ac.jp`

Shiga University

Article

Keywords: scaffold-preserving, lead optimization, structure–activity relationship model, molecular editing

Posted Date: October 5th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11008353/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Molecular descriptors are frequently used to interpret quantitative structure–activity relationship models and to guide lead optimization. However, translating a specified descriptor change into concrete chemical structures remains challenging. Here, we developed Scaffold-tuner, a property-directed compound editing tool that proposes scaffold-preserving structural modifications toward specified changes in hydrogen-bond donors, hydrogen-bond acceptors, aromatic rings, and rotatable bonds. Scaffold-tuner was evaluated using 1,000 drug-like compounds from ChEMBL and compared with random local editing, chemically reasonable mutations (CReM), and REINVENT4 Mol2Mol. Scaffold-tuner achieved the intended descriptor change in seven of eight evaluated tasks and showed particularly high efficiency when only a small number of candidates were considered. With one evaluated candidate, the mean parent success rate was 42.3%, compared with 13.9%, 15.6%, and 17.1% for random editing, CReM, and REINVENT4, respectively. Although broader generative approaches achieved higher coverage with larger candidate budgets, Scaffold-tuner retained relatively high similarity to parent structures. These results suggest that Scaffold-tuner provides a useful framework for translating specified descriptor changes into a limited number of interpretable, parent-like candidate structures.

1. Introduction

In drug discovery, lead optimization aims to improve the physicochemical and pharmacokinetic properties of candidate compounds through iterative structural modification^[1]. Important descriptors of this process include the number of hydrogen-bond donors (HBD), hydrogen-bond acceptors (HBA), aromatic rings (AR), and rotatable bonds (RotB). These descriptors are central to well-known drug-likeness criteria, such as Lipinski's rule of five^[2] and Veber's oral bioavailability rules^[3]. Changes in these descriptors directly affect the molecular properties of the chemical entities being derived. To systematically explore the relationship between these properties and the associated structural changes, a tool that allows the application of well-defined structural interventions to generate candidate compounds is necessary.

In structural-activity relationship studies, examining the effects of defined structural changes can help relate molecular modifications to changes in biological activity or molecular properties^[4, 5, 6]. However, a single structural modification can affect multiple molecular properties simultaneously, which can complicate interpretation of the contribution of an individual descriptor^[5, 6]. Existing approaches such as matched molecular pair analysis (MMPA) primarily derive relationships between structural transformations and observed property changes from existing compound data, rather than explicitly generating structures to achieve a user-specific descriptor change.

MMPA identifies pairs of compounds that differ by a single structural transformation^[4, 6, 7]. MMPA-based tools could suggest R-group replacements based on a historical computational database. However, the suggested transformations are not filtered by their effects on the individual descriptors. A lot of descriptors are related to each other, therefore, multiple descriptors might change simultaneously.

Generative artificial intelligence (AI) tools, such as REINVENT4^[8] and other scaffold-based deep learning methods^[9, 10], could generate numerous novel molecules. However, it is difficult for these tools to generate molecules without constraining the scaffold structure or guaranteeing core scaffold preservation^[9, 11]. Furthermore, interpreting the relationship between the parent and generated compounds is difficult.

Fragment-based drug design^[1] has positioned systematic exploration of chemical space through defined structural modifications as an effective strategy for lead optimization. Based on this concept, in this study, I developed a property-directed compound editing tool “Scaffold-tuner”, which is used for control and interpretable intervention in the elucidation of the scaffold structure of drug-like compounds. This tool uses fragment pools specifically designed to modify each target descriptor. By combining these pools with Murcko decomposition^[12], Scaffold-tuner generates compounds intervened with single structural descriptors while preserving the core scaffold. Based on this feature, compounds having the required number of descriptors are filtered and retrieved.

To confirm the characteristics of Scaffold-tuner, three other algorithms, random local editing, chemically reasonable mutations (CReM)^[13], and Mol2Mol in REINVENT4^[8], were compared using 1,000 compounds collected from ChEMBL, an open database. The methods were evaluated in terms of target descriptor success, candidate efficiency, structural similarity to the parent compounds, scaffold retention, and descriptor selectivity.

2. Results

2.1 Total success ratio

One hundred compounds were generated for each parent compound (collected from the open database) using the simplified molecule input line entry system (SMILES) in four algorithms. The mean and standard deviation of the evaluation scores were calculated as summarized in Table 1.

Table 1

Summary of compound-generation performance across the four methods. Applicability represents the percentage of the 1,000 parent compounds for which at least one valid candidate was generated. Median similarity represents the median Tanimoto similarity between generated candidates and their corresponding parent compounds, calculated using Morgan fingerprints (radius = 2, 2048 bits). Murcko retention represents the percentage of generated candidates retaining the Murcko scaffold of the corresponding parent compound. Δ QED and Δ SA represent the differences in QED and SA Score, respectively, between each generated candidate and its parent compound. Median values are reported for continuous metrics, whereas percentages are reported for applicability and Murcko scaffold retention.

Metric	Scaffold-tuner (proposed)	Random local editing	CReM	Mol2Mol in REINVENT4
Applicability	100%	100%	100%	92.9%
Median similarity	0.721	0.717	0.750	0.520
Murcko retention	75.1%	78%	86.6%	33.6%
Median Δ QED	-0.055	-0.025	-0.026	0.000
Median Δ SA	+ 0.0110	+ 0.270	+ 0.084	+ 0.004

Importing SMILES, Mol2Mol accepted 929 of the 1,000 parent molecules (92.9%) although the other algorithms imported the SMILES of all parent compounds (100.0%). In this study, structures such as [S+] were rejected in REINVENT4. This may be a limitation of using SMILES tokens. Compounds were not generated when the input SMILES contained tokens that were not supported^[8]. This proves that using rule-based algorithms, such as Scaffold-tuner, instead of deep learning-based models is more effective to generate results for variable parent compounds.

2.2 Candidate efficiency under limited candidate budgets

The success ratio of modifying the descriptor values toward the required directions are summarized in Table 2. Scaffold-tuner could generate compounds from parent compounds in 7 out of 8 tasks. In this result, the modified parent compounds achieved a success rate of 48–68% in 6 tasks. Although the proposed tool did not consistently outperform the other algorithms, it demonstrated reasonable level of performance for compound transformation. The performance in HBD + 1, HBA + 1, and AR + 1 was high, whereas the performance in RotB + 1 and AR-1 was low.

Table 2

Parent success rates for the eight descriptor-modification tasks.

Parent success rate was defined as the percentage of parent compounds for which at least one generated candidate achieved the exact target descriptor change of + 1 or - 1. The primary cross-method comparison was performed using the common set of 929 parent compounds accepted by all four methods. Up to 100 candidates per parent were considered for each task. Values are reported as percentages of successful parent compounds; no averaging across repeated runs was performed for this analysis.

Task	Scaffold-tuner (proposed)	Random	CReM	Mol2Mol in REINVENT4
HBD + 1	69.8%	100.0%	69.2%	94.3%
HBD - 1	49.0%	73.3%	43.7%	73.2%
HBA + 1	69.8%	100.0%	72.2%	97.8%
HBA - 1	59.8%	85.6%	43.2%	88.7%
RotB + 1	15.2%	84.8%	83.7%	98.1%
RotB - 1	51.3%	53.3%	26.8%	84.5%
AR + 1	69.8%	1.1%	0.0%	94.7%
AR -1	0.0%	62.9%	0.0%	53.9%

REINVENT4 showed higher parent compound success rates for most tasks when up to 100 candidates were considered, indicating broader coverage of the evaluated descriptor changes. Random local editing also achieved high success rates for several HBD and HBA tasks, whereas CReM showed a more limited task coverage. Notably, Scaffold-tuner achieved a high success rate for AR + 1 (69.8%), while this transformation was rarely observed with random local editing (1.1%) or CReM (0.0%). This result indicates that property-directed fragment selection can enrich specific transformations that are uncommon in descriptor-agnostic local editing.

Because high parent compound success at a large candidate budget does not necessarily indicate efficient proposal of useful candidates, the methods under limited candidate budgets were evaluated next.

2.3 Evaluation under limited candidate repeats

To evaluate candidate efficiency, parent compound success rates were recalculated by limiting the number of evaluated candidates per parent to 1, 5, 10, 20, 50, or 100 (Fig. 1, details are shown in Table S2 in Supporting Information).

At $k = 1$, Scaffold-tuner achieved 42.3%, compared with 13.8% for random local editing, 15.6% for CReM, and 17.1% for REINVENT4 Mol2Mol. The success rate of Scaffold-tuner increased to 47.5% at $k = 5$ and

reached approximately 48% after that.

Random local editing continued to improve as more candidates were considered and surpassed Scaffold-tuner when k exceeded a value of 10. In contrast, CReM only achieved a success rate of 42.6% even at $k = 100$. CReM generated structurally conservative local modifications, showing slightly higher parent similarity than Scaffold-tuner (medium similarity in Table 1). However, its success rate was substantially lower when only a small number of candidates were evaluated. These results indicate that Scaffold-tuner is more effective when a smaller number of property-directed proposals are required. In contrast, broader descriptor-agnostic exploration can achieve higher coverage when a larger candidate limitation is allowed.

Scaffold-tuner is thus capable of generating useful candidates early, although fewer candidates were generated compared to the other algorithms.

2.4 Locality and selectivity of successful edits

To evaluate whether successful descriptor modifications were achieved while preserving structural similarity to the parent compound, the similarity and Murcko scaffold retention of successful candidates were examined. Successful candidates were defined as compounds achieving the exact target descriptor change of ± 1 .

Among successful edits, Scaffold-tuner retained candidates with relatively high similarity to the parent compounds, with median Tanimoto similarities of approximately 0.67–0.75 across the evaluated tasks (Fig. 2). CReM showed a level of parent compound similarity that was equal to or slightly higher than Scaffold-tuner, consistent with its chemically conservative local mutation strategy. In contrast, successful compounds generated by REINVENT4 showed substantially lower similarity to the corresponding parent compounds, with median values of approximately 0.50–0.53. A similar tendency was observed for scaffold preservation (Table S4).

Descriptor selectivity was evaluated using a stricter endpoint requiring an exact ± 1 change in the target descriptor with no changes in the other three evaluated descriptors. Selective success was substantially lower than target success for most tasks (Table S5a), indicating that structural modification of one descriptor was often accompanied by changes in other descriptors. The extent of these collateral changes was strongly task dependent (Table S5b). For example, although RotB increase showed relatively low target success for Scaffold-tuner, 96.1% of its successful candidates exhibited no collateral changes in the other three descriptors. In contrast, AR increase showed high target success but was frequently accompanied descriptor changes.

These results indicate that target success, structural locality, and descriptor selectivity represent distinct aspects of molecular-editing performance.

2.5 Representative structural comparison of molecular edits across methods

Representative molecular edits were compared across the four methods using the same parent compounds and target descriptor tasks. For each example, one candidate was selected from each method according to a predefined rule, and the corresponding changes in HBD, HBA, RotB, AR, and similarity to the parent compound were displayed together with the structures (Fig. 3).

The representative examples reflected the distinct characteristics observed in the quantitative benchmark. In the examples provided, Scaffold-tuner produced localized structural modifications corresponding to the specified descriptor direction while retaining substantial similarity to the parent compounds. Random local editing also generated local modifications, however, the changes were not explicitly directed toward the target descriptor. CReM produced chemically conservative substitutions with high parent similarity^[13], whereas REINVENT4 often yielded candidates matching the evaluated descriptor change through larger structural modifications^[8].

For example, an AR increase illustrated the property-directed behavior of Scaffold-tuner because the corresponding transformation was rarely obtained by random local editing or CReM in the benchmark. In contrast, REINVENT4 readily achieved the same target change but often with lower similarity to the parent structure. These representative structures support the quantitative results by illustrating the trade-off between targeted modification, structural locality, and broader generative exploration.

3. Discussion

In this study, Scaffold-tuner was systematically evaluated against random local editing, CReM, and REINVENT4 Mol2Mol for property-directed compound editing. The benchmark revealed several limitations of Scaffold-tuner, including strong task dependence and limited success for the same transformation modes. Nevertheless, Scaffold-tuner showed competitive performance across multiple evaluation criteria and efficiently enriched candidates exhibiting the intended descriptor change when only a small number of proposals were evaluated.

This characteristic may be particularly useful for connecting quantitative model interpretation with practical molecular design. In QSAR and related analyses, model interpretation can identify molecular descriptors associated with a change in the predicted property and thus provide hypotheses for structural modification^[24]. Translating such descriptor-level interpretations into chemical hypotheses requires methods that can explicitly modify a specified descriptor and propose the corresponding molecular structures. Scaffold-tuner provides a simple, property-directed framework for this purpose, particularly when a limited number of candidate structure are desired.

Future development should focus on expanding the available structural transformations, including broader R-group coverage, to improve both the applicability and efficiency of candidate generation. Such extensions may allow Scaffold-tuner to provide diverse structural proposals while retaining its property-directed and scaffold-preserving design concept.

4. Methods

4.1 Dataset

A dataset of 1,000 drug-like compounds was randomly retrieved from the ChEMBL database ^[14, 15], applying the following filters: molecular type = “small molecule”; molecular weight (MW) = 100–500 Da; logP < = 5; HBD < = 10; HBA < = 20; RB < = 10. Canonical SMILES were generated using RDKit ^[16] and duplicate structures were removed. The 1,000 compounds retained after these processes were used as parent compounds. The summary of physicochemical properties of parent compounds in the dataset is shown in Table 3 and the dataset is provided as supplementary material.

Table 3
Physicochemical profile of the ChEMBL benchmark dataset (n = 1,000)

Descriptor	Mean $\pm$ SD	Min	Max
MW (Da)	327.09 $\pm$ 83.10	106.17	601.56
logP	2.55 $\pm$ 1.68	-9.56	5.00
HBD	1.63 $\pm$ 1.29	0	7
HBA	4.19 $\pm$ 1.76	0	11
RB	4.22 $\pm$ 2.48	0	10
AR	2.00 $\pm$ 1.07	0	5

4.2 Scaffold-tuner (proposed)

Scaffold-tuner is an editing tool used to suggest compounds that are modified with the required descriptors while retaining the core scaffold. If the transformation mode is selected, compounds with the required change are filtered and proposed by referring to a fragment pool based on descriptor definitions and established medicinal chemical data (Table 4). The fragments in each pool were selected based on two criteria: (1) the fragment was expected to modulate the target descriptor based on established medicinal chemical data and descriptor definitions ^[3, 7, 17, 18, 19–21], and (2) the fragments in the pool were required to be structurally diverse, to provide multiple alternatives for structural modification.

Table 4
Fragment pools for all eight transformation modes with design rationale.

Mode	Fragment examples	Design rationale	Reference
add_hbd	-NH ₂ , -NHCH ₃ , -OH, 3-pyrrolyl, indolyl	Introduces N-H or O-H hydrogen-bond donor groups.	[18, 20]
remove_hbd	-CH ₃ , -C ₂ H ₅ , -F, -Cl, -iPr	Replaces HBD-containing groups with non-donor substituents.	[18, 20]
add_hba	-OCH ₃ , -CN, -C(=O)CH ₃ , -N(CH ₃) ₂ , -S(=O) ₂ CH ₃	Introduces heteroatoms with hydrogen-bond acceptor functionality.	[4, 20]
remove_hba	-CH ₃ , -C ₂ H ₅ , -F, -Cl, -Ph	Replaces HBA-containing groups with non- or weakly accepting substituents.	[20, 21]
add_ar	Ph, pyridyl, 2-furyl, 2-thienyl	Introduces an additional aromatic ring.	[19, 20]
remove_ar	-CH ₃ , -C ₂ H ₅ , -nPr, -iPr, -Cy	Replaces aromatic groups with non-aromatic substituents.	[19, 22]
add_rb	-C ₂ H ₅ , -nPr, -nBu, -CH ₂ OCH ₃ , -CH ₂ C(=O)CH ₃	Introduces flexible side chains with additional rotatable bonds.	[3, 20]
remove_rb	-CH ₃ , -F, -Cl, -cPr, -Cy	Replaces flexible groups with small or cyclic substituents.	[3, 22]

Since each descriptor, such as HBA and HBD, may be related, it is difficult to precisely collect only single-descriptor modified compounds. Compounds that are distinctly single-descriptor modified or not may be selected as per the user requirements.

4.3 Random local editing

As a simple base line, molecular structures were randomly modified without using any target molecular properties. For each parent molecule, one of three editing operations was randomly applied: atom replacement, addition of terminal atom, or removal of terminal atom. For atom replacement and terminal-atom addition, an element was randomly selected from C, N, O, F, S, and Cl, which correspond to the major element types used in the fragment templates of Scaffold-tuner. Compounds that failed RDKit sanitization were discarded, and duplicate products generated from the same parent were removed.

4.4 CReM

CReM is a design framework used to generate molecules by replacing a fragment of the input molecule with alternative fragments observed in the same local chemical context from a database of known compounds^[13]. In this study, fragment mutations were generated using `mutate_mol()` function with a ChEMBL-derived CReM fragment database (`chembl33_sa2_f5.db`). Fragment replacement was performed without using a target molecular descriptor. The maximum fragment size parameter was set to `max_size = 1` to restrict mutations to relatively local structural modifications.

4.5 REINVENT4 mol2mol

REINVENT4 mol2mol is a deep learning-based molecular-editing algorithm^[8]. Candidate compounds were generated with the pretrained `mol2mol_medium_similarity` from the prior sampling mode. For each parent compound, 100 compounds were generated using multinomial sampling with a temperature of 1.0 and random seed of 42. Generation was performed without using target descriptors. Parent compounds containing SMILES tokens unsupported by the pretrained model vocabulary were recorded as non-applicable cases and excluded from sampling.

4.6 Evaluation

For each compound in the dataset, up to 100 candidate compounds were evaluated. Invalid structures were discarded after RDKit sanitization, and duplicate products generated from the same parent were removed. Molecular descriptors (the number of HBD, HBA, RotB and AR) were calculated using RDKit, and changes in these descriptors were evaluated relative to the corresponding parent compounds. To evaluate the result of these four algorithms, the following metrics were calculated and compared.

1. Parent success rate: the percentage of parent compounds for which at least one candidate achieved the exact target descriptor change of ± 1 .
2. Limited-budget proposal efficiency: the parent success rate when only the first k candidates per parent were evaluated ($k = 1, 5, 10, 20, 50$, or 100). Candidates were randomly ordered within each parent compound, and the analysis was repeated using multiple random seeds.
3. Selective success ratio: the percentage of parent compounds for which at least one candidate achieved the exact target descriptor change of ± 1 while the other three evaluated descriptors remained unchanged.
4. Similarity to parent: Tanimoto index calculated using morgan fingerprint^[23] (radius = 2, bits = 2048) between the parent and generated compounds. Structural similarity among successful edits was evaluated using candidates that achieved the exact target descriptor change.
5. Murcko scaffold retention: the percentage of generated compounds retaining the Murcko scaffold of the corresponding parent compound.

Using these comparative criteria, Scaffold-tuner was evaluated as a competent scaffold editing tool to interpret quantitative structure–activity relationship models and to guide lead optimization.

Declarations

Acknowledgements

I would like to thank Editage (www.editage.jp) for the English language editing.

Funding

This study was supported by the Japan Society for the Promotion of Science, Grants-in-Aid for Scientific Research [grant number 26K09775] and grant for Special Research Promotion Project in Shiga University.

Author Contributions

TE acquired the funding, developed the editing tool, investigated the experimental methods, and formally analysed and validated the experimental data. TE prepared all tables and figures made available in the text and wrote and reviewed the complete content of this manuscript.

Competing Interests

The author declares no competing interests.

Data Availability

The ChEMBL compound dataset used in the benchmark study ($n = 1,000$) is provided as supplementary material. The source code of the Scaffold-tuner is freely available at <https://github.com/tesaki2019/Scaffold-tuner/> under the MIT license. The notebooks, configuration file, and analysis functions used to reproduce the benchmark experiments are available at <https://github.com/tesaki2019/Scaffold-tuner-experiments/>.

References

1. Murray, C. W. & Rees, D. C. The rise of fragment-based drug discovery. *Nat. Chem.* **1** (3), 187–192. 10.1038/nchem.217 (2009).
2. Lipinski, C. A., Lombardo, F., Dominy, B. W. & Feeney, P. J. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* **46** (1), 3–26. [https://doi.org/10.1016/S0169-409X\(00\)00129-0](https://doi.org/10.1016/S0169-409X(00)00129-0) (2001).
3. Veber, D. F. et al. Molecular Properties That Influence the Oral Bioavailability of Drug Candidates. *J. Med. Chem.* **45** (12), 2615–2623. 10.1021/jm020017n (2002).
4. Yang, Z. et al. Matched Molecular Pair Analysis in Drug Discovery: Methods and Recent Applications. *J. Med. Chem.* **66** (7), 4361–4377. 10.1021/acs.jmedchem.2c01787 (2023).
5. Leach, A. G. et al. Matched Molecular Pairs as a Guide in the Optimization of Pharmaceutical Properties; a Study of Aqueous Solubility, Plasma Protein Binding and Oral Exposure. *J. Med. Chem.*

- 49** (23), 6672–6682. 10.1021/jm0605233 (2006).
6. Dossetter, A. G., Griffen, E. J. & Leach, A. G. Matched Molecular Pair Analysis in drug discovery. *Drug Discovery Today*. **18** (15), 724–731. <https://doi.org/10.1016/j.drudis.2013.03.003> (2013).
 7. Hussain, J. & Rea, C. Computationally Efficient Algorithm to Identify Matched Molecular Pairs (MMPs) in Large Data Sets. *J. Chem. Inf. Model.* **50** (3), 339–348. 10.1021/ci900450m (2010).
 8. Loeffler, H. H. et al. Reinvent 4: Modern AI-driven generative molecule design. *J. Cheminform.* **16** (1), 20. 10.1186/s13321-024-00812-5 (2024).
 9. Li, Y. et al. A Comprehensive Tool for Scaffold-Based De Novo Drug Discovery Using Deep Learning. *J. Chem. Inf. Model.* **60** (1), 77–91. 10.1021/acs.jcim.9b00727 (2019).
 10. Langevin, M., Minoux, H., Levesque, M. & Bianciotto, M. Scaffold-Constrained Molecular Generation. *J. Chem. Inf. Model.* **60** (12), 5637–5646. 10.1021/acs.jcim.0c01015 (2020).
 11. Hu, Y. & Bajorath, J. r. Exploring Target-Selectivity Patterns of Molecular Scaffolds. *ACS Med. Chem. Lett.* **1** (2), 54–58. 10.1021/ml900024v (2010).
 12. Bemis, G. W. & Murcko, M. A. The Properties of Known Drugs. 1. Molecular Frameworks. *J. Med. Chem.* **39** (15), 2887–2893. 10.1021/jm9602928 (1996).
 13. Polishchuk, P. CReM: chemically reasonable mutations framework for structure generation. *J. Cheminform.* **12** (1), 28. 10.1186/s13321-020-00431-w (2020).
 14. Bento, A. P. et al. An open source chemical structure curation pipeline using RDKit. *J. Cheminform.* **12** (1), 51. 10.1186/s13321-020-00456-1 (2020).
 15. Zdrazil, B. Fifteen years of ChEMBL and its role in cheminformatics and drug discovery. *J. Cheminform.* **17** (1). 10.1186/s13321-025-00963-z (2025).
 16. *rdkit/rdkit: 2026_03_6 (Q1 2026) Release (Versions Release_2026_03_6)*; Zenodo: (accessed. 9/11/2026) (2026).
 17. Takeuchi, K., Kunimoto, R. & Bajorath, J. Searchable database of frequent R-groups in medicinal chemistry and their preferred replacements. *Data Brief.* **39**, 107456. 10.1016/j.dib.2021.107456 (2021).
 18. Kenny, P. W. Hydrogen-Bond Donors in Drug Design. *J. Med. Chem.* **65** (21), 14261–14275. 10.1021/acs.jmedchem.2c01147 (2022).
 19. Ward, S. E. & Beswick, P. What does the aromatic ring number mean for drug design? *Expert Opin. Drug Discov.* **9** (9), 995–1003. 10.1517/17460441.2014.932346 (2014).
 20. Takeuchi, K., Kunimoto, R. & Bajorath, J. Systematic mapping of R-group space enables the generation of an R-group replacement system for medicinal chemistry. *Eur. J. Med. Chem.* **225**, 113771. <https://doi.org/10.1016/j.ejmech.2021.113771> (2021).
 21. Pollock, J. et al. Rational Design of Orthogonal Multipolar Interactions with Fluorine in Protein–Ligand Complexes. *J. Med. Chem.* **58** (18), 7465–7474. 10.1021/acs.jmedchem.5b00975 (2015).
 22. Lovering, F., Bikker, J. & Humblet, C. Escape from Flatland: Increasing Saturation as an Approach to Improving Clinical Success. *J. Med. Chem.* **52** (21), 6752–6756. 10.1021/jm901241e (2009).

23. Rogers, D., Hahn, M., Extended-Connectivity & Fingerprints *J. Chem. Inf. Model.* **50** (5), 742–754. DOI: 10.1021/ci100050t. (2010).
24. Polishchuk, P. Interpretation of Quantitative Structure–Activity Relationship Models: Past, Present, and Future. *J. Chem. Inf. Model.* **57** (11), 2618–2639. 10.1021/acs.jcim.7b00274 (2017).

Figures




Figure 1

Parent compound success rate as a function of the number of evaluated candidates. Parent compound success rates were calculated for each method by limiting the number of evaluated candidates per parent and task to 1, 5, 10, 20, 50, or 100. For each candidate budget, candidates were randomly sampled 20 times, and the mean parent compound success rate was calculated across the eight descriptor-modification tasks. Error bars indicate the standard deviation across the repeated samplings.




Figure 2

Structural similarity to parent compounds among successful descriptor modification. Median Tanimoto similarity to the corresponding parent compound was calculated for candidates achieving the exact target descriptor change of ± 1 in each task. Tanimoto similarity was calculated using Morgan fingerprints (radius=2, 2048 bits). Higher values indicate greater structural similarity to the parent compound.




Figure 3

Representative compounds. Representative examples were selected using a predefined procedure and were not selected based on visual preference of Scaffold-tuner performance.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [chembl1000parents.csv](#)
- [ScaffoldtunerEsakiSR2026SI.docx](#)