

Article

Do Molecular Fingerprints Identify Diverse Active Drugs in Large-Scale Virtual Screening? (No)

Vishwesh Venkatraman ^{1,*} , Jeremiah Gaiser ², Daphne Demekas ³ , Amitava Roy ^{4,5} , Rui Xiong ⁶ 
and Travis J. Wheeler ^{3,*} 

¹ Department of Chemistry, Norwegian University of Science and Technology, 7034 Trondheim, Norway

² School of Information, University of Arizona, Tucson, AZ 85721, USA

³ R. Ken Coit College Pharmacy, University of Arizona, Tucson, AZ 85721, USA

⁴ Rocky Mountain Laboratories, Bioinformatics and Computational Biosciences Branch, Office of Cyber Infrastructure and Computational Biology, National Institute of Allergy and Infectious Diseases, National Institutes of Health, Hamilton, MT 59840, USA; amitava.roy@umontana.edu

⁵ Department of Biomedical and Pharmaceutical Sciences, University of Montana, Missoula, MT 59812, USA

⁶ Department of Pharmacology & Toxicology, University of Arizona, Tucson, AZ 85721, USA

* Correspondence: vishwesh.venkatraman@ntnu.no (V.V.); twheeler@arizona.edu (T.J.W.)

Abstract: Computational approaches for small-molecule drug discovery now regularly scale to the consideration of libraries containing billions of candidate small molecules. One promising approach to increased the speed of evaluating billion-molecule libraries is to develop succinct representations of each molecule that enable the rapid identification of molecules with similar properties. Molecular fingerprints are thought to provide a mechanism for producing such representations. Here, we explore the utility of commonly used fingerprints in the context of predicting similar molecular activity. We show that fingerprint similarity provides little discriminative power between active and inactive molecules for a target protein based on a known active—while they may sometimes provide some enrichment for active molecules in a drug screen, a screened data set will still be dominated by inactive molecules. We also demonstrate that high-similarity actives appear to share a scaffold with the query active, meaning that they could more easily be identified by structural enumeration. Furthermore, even when limited to only active molecules, fingerprint similarity values do not correlate with compound potency. In sum, these results highlight the need for a new wave of molecular representations that will improve the capacity to detect biologically active molecules based on their similarity to other such molecules.

Keywords: molecular fingerprints; drug screening; similarity searching



Citation: Venkatraman, V.; Gaiser, J.; Demekas, D.; Roy, A.; Xiong, R.; Wheeler, T.J. Do Molecular Fingerprints Identify Diverse Active Drugs in Large-Scale Virtual Screening? (No). *Pharmaceuticals* **2024**, *17*, 992. <https://doi.org/10.3390/ph17080992>

Academic Editors: George Mihai Nitulescu, Roberta Rocca and Dragos Mihai

Received: 8 June 2024

Revised: 18 July 2024

Accepted: 23 July 2024

Published: 26 July 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Methods for the computational identification of small molecules likely to bind to a drug target (virtual screening) are increasingly intended to explore a space of billions of candidate molecules [1–4]. One strategy for exploring this massive molecular search space is to begin with a collection of molecules known or presumed to be biologically active (hereafter referred to as *actives*). Those actives can then be used as the basis of a rapid search among billions of candidates [5] for other molecules expected to demonstrate similar activity [6,7].

The notion that small molecules with similar structures are likely to share biological properties, coined the *similar property principle* (SPP) [8,9], is central to this sort of search strategy. The SPP is simple and intuitive, and has served as the basis for predictions of biological activity [10,11], toxicity [12–14], aqueous solubility [14,15] (logS), and partition coefficient [16] (logP). However, structural similarity may not necessarily reflect similarity in biological activity [17], a concept popularly addressed by the notion of an Activity Cliff [18–20]. Furthermore, the proper definition of structural similarity depends on the

context. For example, the quantitative structure–activity relationship focuses on the similarity between the local structural features of two molecules, while similarity in biological activity typically depends on having more global features of the molecules [21–23] (though even these notions of *local* and *global* similarity are also not well defined).

The most common way to quantify the structural similarity of two small molecules begins with the calculation of a so-called *molecular fingerprint*, a binary or count vector that encodes structural and often chemical features [24–26]. Such a fingerprint is computed for each molecule, then the fingerprints of molecules are compared for overlap to approximately assess molecular similarity, motivated by the notion that this measure of fingerprint similarity will correlate with the similarity at the level of biological activity. Fingerprint similarity has been used to effectively estimate $\log S$ and $\log P$ values [27], likely because these values can largely be approximated from the small molecule itself without explicitly considering interacting partners.

Other molecular properties involve a dependency on context, placing greater strain on the utility of the SPP. For example, the biological activity of a small molecule depends on the interaction between that molecule and the target protein binding region. Such binding regions (or pockets) vary between different proteins, and therefore impose strong context dependence in biological interactions. Consequently, small molecule ligand-based fingerprint similarity may not be sufficient to capture the wide spectrum of similarities in biological activities. Similarly, the toxicity of a small molecule depends on the molecule's interaction with multiple proteins, suggesting a limit to the inference power provided by similarity at the level of molecular fingerprints.

Despite previous demonstrations of the limitations of using SPP to predict similarity in biological activity [28,29], the technique is heavily used in drug development. This is especially true in fingerprint-based virtual screening (VS), in part due to the computational simplicity and speed of searching the vast chemical space of small molecules [2,30–32]. A variety of molecular fingerprints have been devised for use in ligand-based virtual screening (LBVS), to aid in identifying biologically active small molecules [32,33] (hereafter referred to as actives), from within a library of small molecules. LBVS begins with a small number of known/predicted actives (queries) for a target protein pocket, and explores a library of other small molecules, seeking a set that is expected to also be active. This expectation is based on the SPP, so that LBVS seeks molecules with high fingerprint similarity to one of the queries, under the assumption that fingerprint similarity to known actives will generally assign a higher ranking to actives than to non-actives (decoys). Martin et al. [28] have highlighted a disconnect between empirical and computational perception of similarity, and suggested resulting limitations in the utility of LBVS for drug binding predictions. Here, we explore the shortcomings of simplistic molecular fingerprints in the context of LBVS, and demonstrate that all commonly used fingerprint methods fail to sufficiently enrich for binding activity in a library of mostly decoy molecules. We are not oblivious to a prevailing perception that fingerprint similarity is suggestive of similarity in drug interaction potential, and hope that this manuscript contributes to a better understanding of the limitations of fingerprint similarity in the context of large-scale drug screening.

2. Results

To gain insight into the utility of various fingerprinting strategies for billion-scale virtual drug screening, we explored the capacity of fingerprint similarity to extract a small set of candidates that is highly enriched for molecules with activity similar to the seed query molecules. First, we computed measures of enrichment for 32 fingerprints on four benchmark data sets, presenting both classical ROC AUC calculations and our new decoy retention factor (DRF) scores. We then explored the distributions of fingerprint similarity scores across a variety of target molecules, and show that the score distributions for actives and decoys are not sufficiently separable for billion-scale search. We also found that, where fingerprints did occasionally find good matches among actives, the “new” actives almost always shared a scaffold with the query active (meaning that they would be found by

more traditional shape exploration methods). We further considered whether there is a correlation between compound potency and active–active similarity scores, and found that there is not. Finally, we used a data set containing more than 300,000 experimentally confirmed inactive compounds, and found that fingerprint similarity to an active molecule does not enable discrimination between actives and inactive. In total, these results indicate that fingerprint similarity is not a reliable proxy for likely similar binding activity or particularly for discovering diverse active structures.

2.1. Enrichment for Active Molecules

To assess the utility of fingerprinting strategies for selecting compounds with similar expected activity, we computed the similarities of all compounds to a query active molecule, and tested whether active molecules tend to be more similar to other actives than to decoys. Specifically, for each target protein, we computed the fingerprints of each molecule associated with that target protein. Then, for each active compound, we computed the similarity of its fingerprint to each of the other compounds (actives and decoys) affiliated with that target. The union of these active/decoy distance calculations was merged and sorted by similarity, enabling calculation of $DRF_{0.1}$ (see Section 4) and ROC AUC for each fingerprint–target pair. Then, for each fingerprint–target pair, the mean value (AUC or $DRF_{0.1}$) was computed over all target proteins in the corresponding benchmark.

Table 1 presents the resulting enrichment values on each benchmark data set. The performance of all fingerprints is poor for both the MUV and LIT-PCBA data sets, with AUC values generally <0.6 , and $DRF_{0.1}$ values close to 1.0 (indicating small enrichment of actives relative to decoys). Performance is somewhat better on DEKOIS and DUD-E, but not particularly strong, and is offset by concerns of benchmark bias highlighted elsewhere, such as artificial enrichment [34–39] (enrichment due to bias in the actives or decoys), analogue bias (limited diversity of the active molecules), and false negative bias (risk of active compounds being present in the decoy set), all of which can cause misleading VS results [37,39,40]. Table 1 also provides a summary of the VS performances obtained for the fingerprint types (substructure, circular, path, text, pharmacophore). No particular fingerprint strategy appears to be better suited to the problem of virtual screening. Among all tests, the best observed $DRF_{0.1}$ value was 0.09, equating to a roughly 11-fold enrichment in actives versus inactives; in a screen of a billion molecules, this means that roughly 9 million inactive molecules are expected to show similarity at the score threshold where 10% of actives are retained).

Most circular and path-based fingerprints employ a standard length of 1024 bits. O’Boyle and Sayle [41] suggested that increasing the bit-vector length from 1024 to 16,384 can improve VS performance, though at a cost of space and run time for comparison. We evaluated the utility of longer fingerprints for the MUV and LIT-PCBA data sets, and found that longer fingerprints yield little to no gain in efficacy (see Table S1 in Supplementary Information).

2.2. Tanimoto Similarity Distributions Are Generally Indistinguishable

To explore the distribution of similarities between actives and decoys, we computed Tanimoto coefficients for active–active and active–decoy molecule pairs in the DEKOIS data set. For each target protein in DEKOIS [42], we randomly selected an active molecule, and computed the molecular fingerprint Tanimoto similarity to all other actives and decoys for that target. Figure 1 shows the resulting score distributions for 32 fingerprints. The distributions of active–active (blue) and active–decoy (red) Tanimoto values substantially overlap, so that the vast majority of actives fall into a score range shared by most decoys. (Note: a single active molecule was used as query in order to replicate a common use case for fingerprints in virtual screening. Similarly overlapping distributions are observed when merging the results of using all DEKOIS actives as query seeds, but with a more complex visual landscape due to a mixture of similarity distributions.)

Table 1. Summary of the VS performances in terms of the AUC and DRF ($p = 0.1$) for the 32 fingerprints tested on the DEKOIS, DUDE, MUV and LIT-PCBA data sets. Each row with no background corresponds to a fingerprint. Each row with a grey background corresponds to the family of fingerprints given in the preceding rows, and shows the mean value for that family. See also heatmaps of the different metrics for the data sets in Figures S1–S5 in the Supplementary Information.

FP	AUC				DRF			
	DEKOIS	DUDE	MUV	LIT-PCBA	DEKOIS	DUDE	MUV	LIT-PCBA
AP2D	0.64	0.66	0.49	0.51	0.55	0.39	1.24	1.00
AT2D	0.78	0.79	0.58	0.55	0.20	0.12	0.86	0.83
AVALON	0.72	0.73	0.60	0.55	0.30	0.18	0.85	0.97
ESTATE	0.71	0.75	0.53	0.50	0.31	0.17	0.98	0.94
FP3	0.68	0.77	0.51	0.52	0.45	0.20	0.89	0.83
FP4	0.74	0.80	0.58	0.54	0.30	0.12	0.89	0.91
MACCS	0.71	0.75	0.55	0.54	0.33	0.18	0.99	0.93
PUBCHEM	0.76	0.76	0.55	0.54	0.28	0.20	1.06	0.97
RDK5	0.76	0.75	0.58	0.56	0.24	0.16	0.90	0.89
RDK6	0.70	0.70	0.59	0.58	0.38	0.29	0.86	0.82
RDK7	0.62	0.63	0.58	0.59	0.74	0.70	0.98	0.85
KR	0.72	0.74	0.54	0.51	0.31	0.19	1.05	0.96
SIGNATURE	0.72	0.72	0.55	0.54	0.33	0.23	0.97	0.99
SUBSTRUCTURE	0.71	0.73	0.56	0.54	0.36	0.24	0.96	0.91
ASP	0.80	0.79	0.58	0.53	0.17	0.12	0.89	0.99
DFS	0.79	0.78	0.55	0.52	0.18	0.14	0.98	1.01
FP2	0.79	0.78	0.55	0.54	0.20	0.14	1.01	0.90
LSTAR	0.78	0.78	0.54	0.51	0.19	0.13	0.97	1.03
TT	0.80	0.80	0.61	0.56	0.15	0.10	0.72	0.85
PATH	0.75	0.75	0.57	0.54	0.18	0.13	0.91	0.96
ECFP0	0.70	0.77	0.53	0.50	0.33	0.13	0.97	0.96
ECFP2	0.77	0.81	0.54	0.51	0.19	0.09	0.99	1.01
ECFP4	0.76	0.80	0.54	0.51	0.19	0.09	0.99	1.00
ECFP6	0.75	0.78	0.54	0.52	0.20	0.10	0.99	0.98
FCFP0	0.66	0.69	0.54	0.52	0.35	0.23	0.41	0.42
FCFP2	0.76	0.75	0.55	0.52	0.24	0.19	0.93	1.01
FCFP4	0.78	0.76	0.54	0.52	0.20	0.15	0.93	1.01
FCFP6	0.78	0.75	0.54	0.52	0.20	0.15	0.96	0.98
MAP4	0.81	0.83	0.56	0.54	0.14	0.07	0.91	0.85
MHFP	0.81	0.81	0.54	0.53	0.17	0.10	0.97	0.94
RAD2D	0.76	0.77	0.53	0.53	0.23	0.14	0.99	0.93
CIRCULAR	0.76	0.77	0.54	0.52	0.26	0.11	0.98	0.99
P2PPHAR2D	0.66	0.74	0.51	0.54	0.50	0.25	1.20	0.94
P3PPHAR2D	0.71	0.76	0.52	0.55	0.33	0.16	1.17	0.93
PHARMACOPHORE	0.75	0.76	0.54	0.53	0.42	0.21	1.19	0.94
LINGO	0.77	0.79	0.54	0.54	0.21	0.10	1.02	0.91

Most of the fingerprints in Figure 1 present a thin high-Tanimoto tail for actives (blue) that is not seen for decoys (red), suggesting that perhaps a small fraction of actives could be discriminated from decoys by establishing a high score threshold—in other words, it may be possible to select a threshold that delineates regions of early enrichment. However, consider the ECFP2 fingerprint, which shows an apparently compelling right tail in the active–active plot (blue), such that it appears to be reasonable to establish a Tanimoto cutoff of 0.5. In DEKOIS, there are 423 active matches to active queries above this threshold. Though the right tail of the active–decoy distribution (red) is imperceptible in this plot, it still contains $\sim 0.0064\%$ of the decoys. Extrapolating to a library of 3.7 billion candidates, as we used in Venkatraman et al. [2], we expect to see $\sim 23.7\text{M}$ decoys with Tanimoto ≥ 0.5 , so that the active-to-decoy ratio is $\sim 1:56,000$. Setting the Tanimoto threshold to 0.75 leads to an expected ratio of $\sim 1:68,000$ (57 actives to $\sim 3.9\text{M}$ expected decoys). This is likely an overly pessimistic view of the decoy ratio risk, since the compounds in the decoy set are intended to be more similar to the actives than would be a random molecule from a billion-scale library. Even so, it highlights the fact that even small false discovery rates

can produce an overwhelmingly large number of false matches when the target base is sufficiently large. Moreover, see the next section for an exploration of the expected novelty of these high-scoring matches.

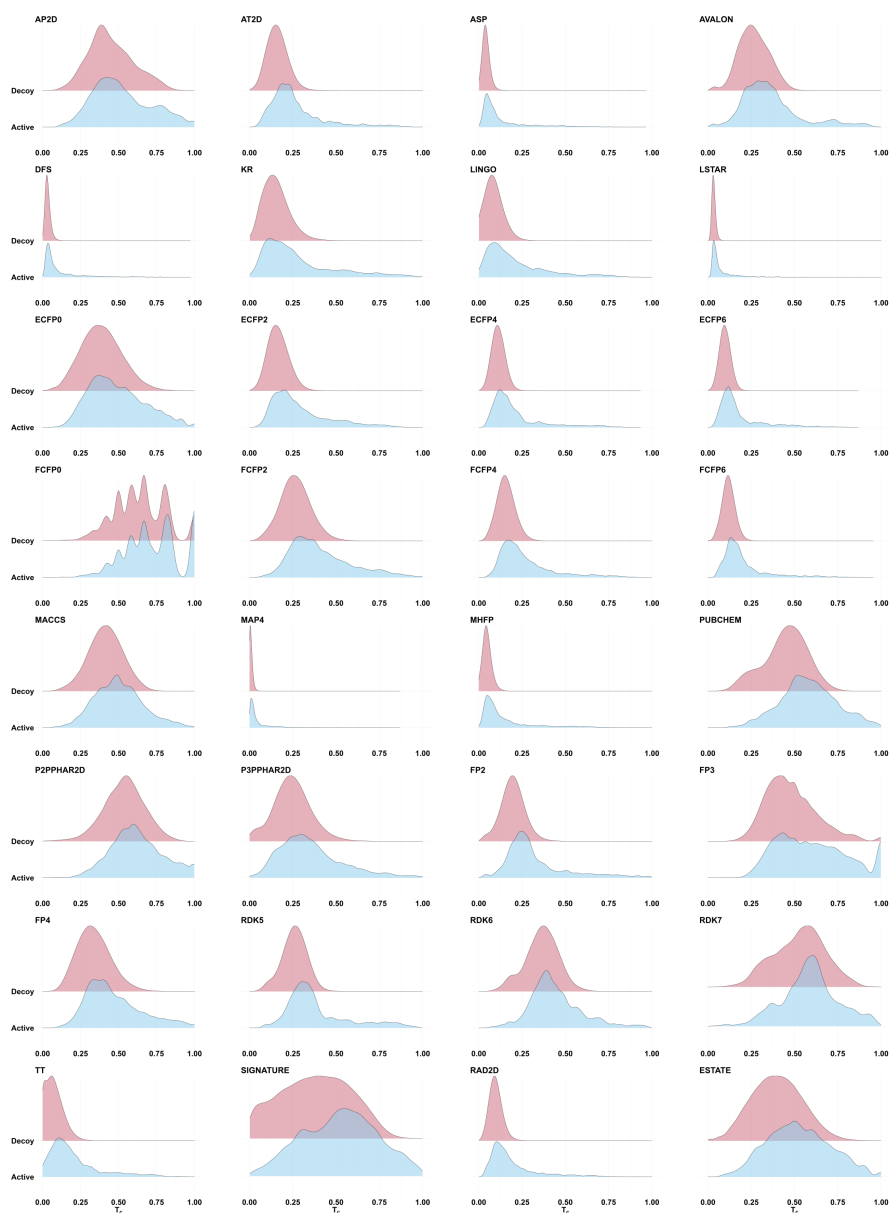


Figure 1. Ridgeline plots showing the distribution of the Tanimoto fingerprint similarities calculated between a randomly selected active molecule for each target protein and all other actives (shown in blue) and decoys (in red) for that target. Data taken from the DEKOIS data set. The distribution of similarity scores between the active query and other active molecules is largely indistinguishable from the distribution of similarity scores to random molecules. Where the active (blue) distribution does show a fatter high-scoring tail than the inactive distribution (suggesting potential for early enrichment by using a high score threshold), a search against a large target database will still produce filtered sets that are massively dominated by inactives (see text).

2.3. Tanimoto Score Distributions, and Their Utility in Scaffold Hopping

The LIT-PCBA data set [43] contains binding activity data for 15 proteins, each with 13 to 1563 confirmed active molecules and 3345 to 358,757 confirmed inactives/decoys (see Table 2). We computed the ECFP4 fingerprint for each of the 2.6 M active and inactive small molecules. For each protein, we first selected the molecule with best affinity to serve as the

‘active query’, then computed the Tanimoto similarity measure (T_c) between seed and all other molecules (actives and decoys) for that protein. Figure 2 presents T_c distributions. Each plot shows, for a single protein, the fraction of actives (blue) and decoys (red) that are recovered (Y-axis) as a function of T_c (X-axis). One important observation is that for most proteins, the molecule with the highest similarity (active or decoy) to the active query has a $T_c < 0.5$, which is below commonly used thresholds for expected shared bioactivity [44–46]. In addition, for most proteins, T_c distributions are indistinguishable, though three proteins (GBA, OPRK1, and PPARG) demonstrate early enrichment for actives. We manually inspected all high-scoring actives ($T_c > 0.5$) for these three enriched proteins (see Figure 3 and Supplementary Figures S6 and S7), and observed that high-scoring actives were almost entirely bioisosteres, with easily predictable results of exchanging an atom or atom-group with a similar atom or group. The lack of scaffold diversity among matches with scores discernible from high-volume decoy noise casts doubt on the utility of fingerprint similarity for exploring novel candidate drug spaces.

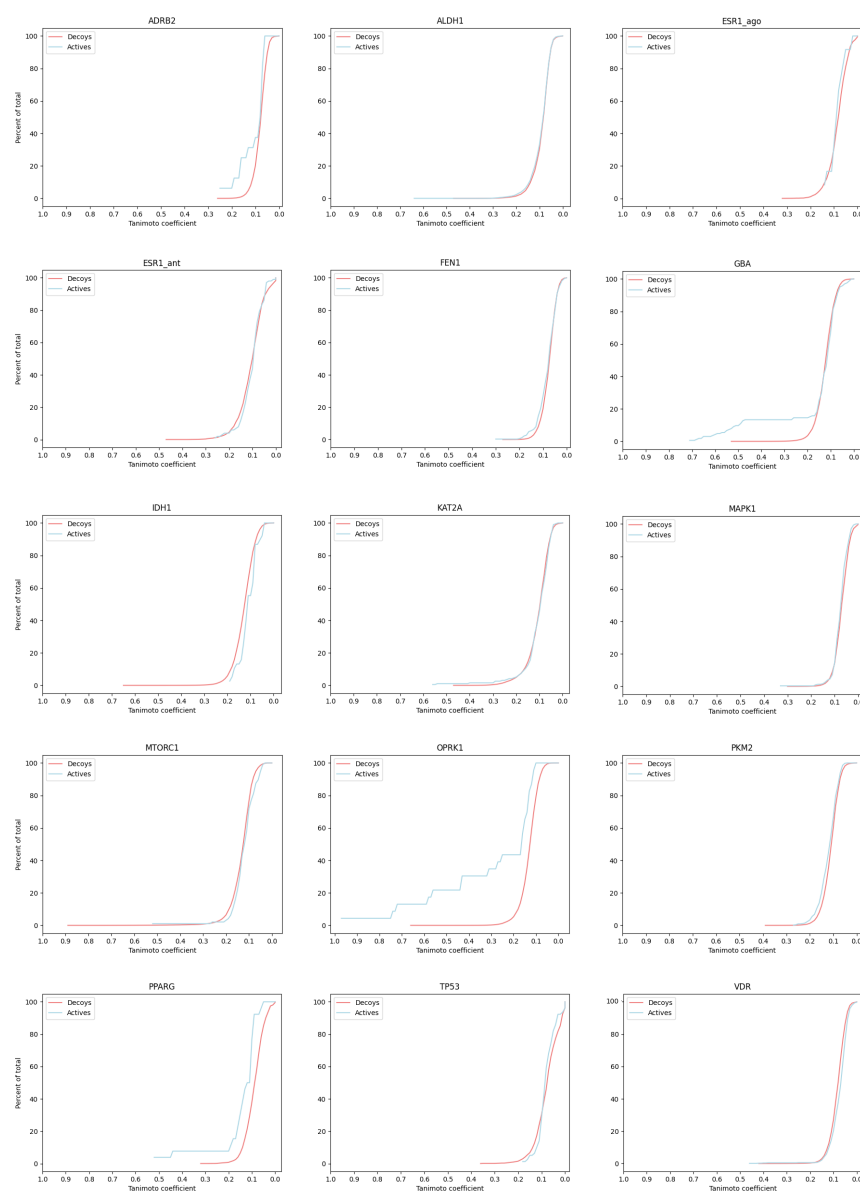








Figure 2. Cumulative percent of total actives (blue) and decoys (red) encountered (Y-axis) as a function of decreasing Tanimoto coefficient, using ECFP4 fingerprints. For each protein, the molecule with the best binding affinity was used as the query molecule (the molecule for which the Tanimoto score was computed for each other molecule, active or inactive).

OPRK1**Active Query**
CCC(C)(C(NC1CCCC1)=O)N(Cc1cccs1)C(c1ncccc1)=O
**Active hits with Tanimoto >=0.5 (ECFP4)**
CCC(C)(C(NC1CCCC1)=O)N(Cc1cccs1)C(c1ncccc1)=O

CCC(C)(C(NC1CCCC1)=O)N(Cc1ccco1)C(c1ncccc1)=O

CCC(C)(C(NC1CCCC1)=O)N(Cc1ccco1)C(c1ncccc1)=O

CCC(C)(C(NC1CCCC1)=O)N(Cc1ccco1)C(c1nccnc1)=O

O=C(C1(CCCCC1)N(Cc1cccs1)C(c1ncccc1)=O)NC1CCCC1


Same scaffold, obvious

Figure 3. In the LIT-PCBA analysis, early enrichment was observed for OPRK1—for all Tanimoto coefficients $t > \sim 0.2$, the fraction of actives with Tanimoto score $> t$ is much larger than the fraction of decoys with that score (see Figure 2). This suggests that fingerprints produce useful early enrichment of active molecules, at least for this one protein target. We sought to understand if the high-scoring actives represented novel drug candidates that could not be easily identified by simple modifications to the active drug used as a fingerprint query. There are only 24 actives in the OPRK1 data set, and only 5 of these showed Tanimoto score > 0.5 to the initial query; we manually inspected the structures of these compounds. All 5 are built on the same scaffold as the query (in red), and are obvious variations that should be identified through standard enumeration (i.e., no new scaffold are explored). Similar plots are provided for GBA and PPARG in Supplementary Figures S6 and S7.

Table 2. LIT-PCBA provides sets of confirmed actives and inactives for 15 protein binding partners.

Protein	Actives	Inactives
ADRB2	17	311,748
ALDH1	5363	101,874
ESR_ago	13	4,378
ESR_antago	88	3,820
FEN1	360	350,718
GBA	163	291,241
IDH1	39	358,757
KAT2A	194	342,729
MAPK1	308	61,567
MTORC1	97	32,972
OPRK1	24	269,475
PKM2	546	244,679
PPARG	24	4,071
TP53	64	3,345
VDR	655	262,648

2.4. Evaluation on a Target with Many Validated Inactive Molecules

The previous experiments depend on benchmarks containing computationally identified decoys that almost entirely have not been experimentally validated as inactive. The MMV St. Jude malaria data set [47] provides an alternative perspective on the utility of fingerprint similarity for activity prediction in the context of verified decoys. It contains a set of 305,810 compounds that were assayed for malaria blood stage inhibitory activity. Among these molecules, 2507 were classified as active, while the remaining 303,303 compounds were classified as inactive.

For each active molecule, we computed Tanimoto similarity to every other active and all inactives. Figure 4 shows bar plots for each fingerprint, with each plot showing the fraction of inactives (red) and actives (blue) with Tanimoto similarity values $T_{ab} \geq c$ for values of $c = (0.1, 0.2, \dots, 0.9)$ and 0.99. In general, the remaining fraction of actives only slightly exceeds the remaining fraction of inactives, indicating minimal enrichment of actives at increased Tanimoto similarity values. MAP4 shows an apparent relative abundance of actives, but note enrichment is still only ~ 10 -fold; also note that $< 1.5\%$ of actives show Tanimoto similarity > 0.1 to another active, raising concerns about how to establish meaningful MAP4 thresholds.

One important caveat for this data set is that the target binding pocket of each active is unknown, so that it is possible that some actives target one pocket while other actives target another pocket. Even so, the lack of visible signal of predictive enrichment from fingerprint similarity is notable.

2.5. Fingerprint Similarity Values Do Not Correlate with Compound Potency

The previous sections demonstrate that fingerprint similarity has limited utility in discriminating active molecules from decoys. An alternative use of fingerprints could be to take a set of candidates that have already (somehow) been highly enriched for active compounds, and rank them according to expected potency. For the LIT-PCBA data set, these potency values for drug–target pairs were retrieved from PubChem [48]: AC_{50} (“Half-maximal Activity Concentration”, defined as the concentration at which the compound exhibits 50% of maximal activity).

For each target protein in the LIT-PCBA data set, we selected the most potent active molecule, and computed fingerprint similarities for all other actives for the corresponding target. We evaluated the correlation of fingerprint similarity value to observed AC_{50} by computing the Kendall rank correlation [49]. Figure 5 presents a heatmap of these correlation values (τ) for each fingerprint across 15 protein targets, and demonstrates that all fingerprints exhibit poor correlation, with values ranging between -0.53 to 0.54 , and generally only slightly higher than zero. This indicates that the fingerprints evaluated are unlikely to yield a ranked set of enriched highly potent compounds, in agreement with the observations of Vogt and Bajorath [50].

To give insight into these summary statistics represented by the heatmap squares in Figure 5, we prepared scatter plots corresponding to three of the target/fingerprint pairings (Figure 6). The middle plot shows the relationship between fingerprint similarity and AC_{50} values for the target/fingerprint pair with median Kendall correlation (target = VDR, fingerprint = SIGNATURE), and is representative of most of the target/fingerprint pairs; it shows essentially no correlation between fingerprint similarity and AC_{50} values ($\tau = 0.01$). The first and last scatter plots show fingerprint similarity and AC_{50} values for the target/fingerprint with the highest (ADRB2, FFCP0) and lowest (PPARG, ASP) correlation values. Note that Kendall rank correlation values for FFCP0 with targets other than ADRB2 ($\tau = 0.54$) vary from -0.30 to 0.05 and for ASP with targets other than PPARG ($\tau = -0.53$) from -0.23 to 0.18 . Though some protein/fingerprint combinations show apparent correlation between fingerprint similarity and AC_{50} , most show a near total lack of correlation; deviations (positive or negative correlation) are observed in proteins with a small number of actives, which are thus more prone to stochastic or biased observations of correlation.

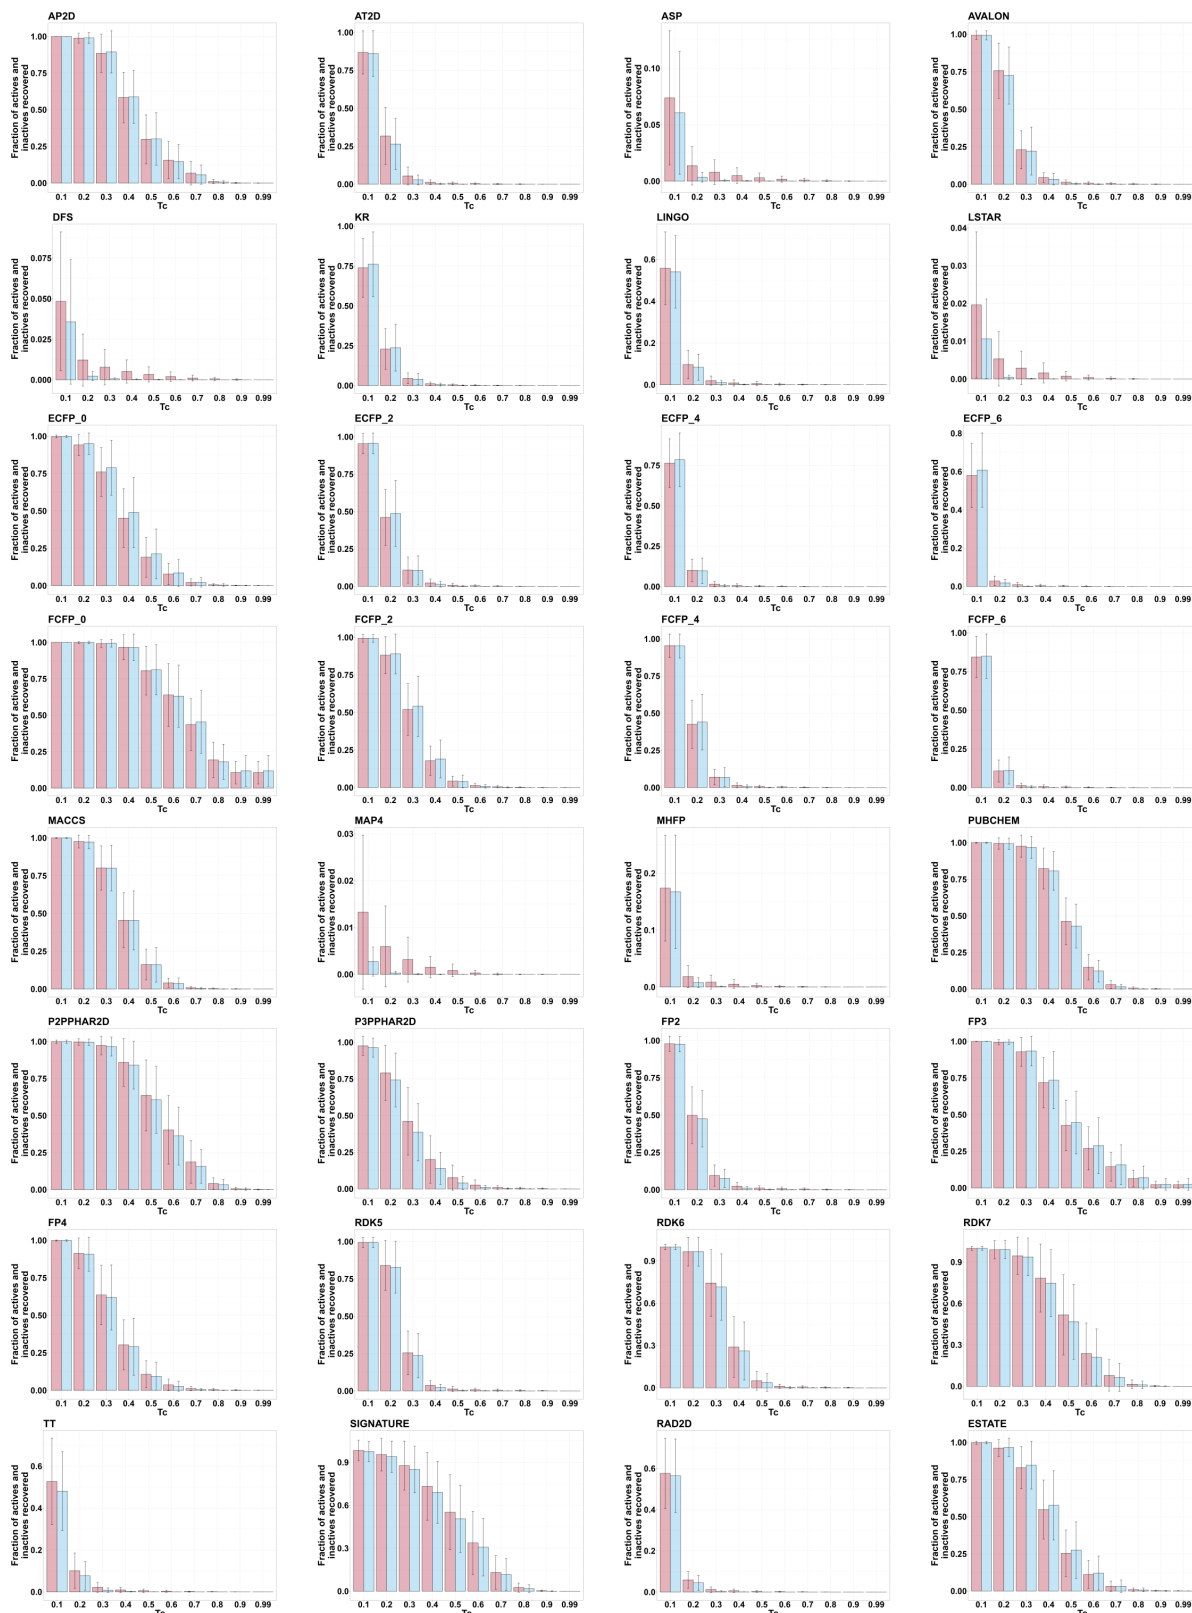


Figure 4. For compounds in the St. Jude malaria data set, bar plots show the fraction of the inactives (in red) and actives (in blue) exceeding Tanimoto similarity cutoffs by the different fingerprints. Tanimoto similarities were calculated using each active as the query; mean and standard deviation (based on the 2507 actives) are shown as error bars.

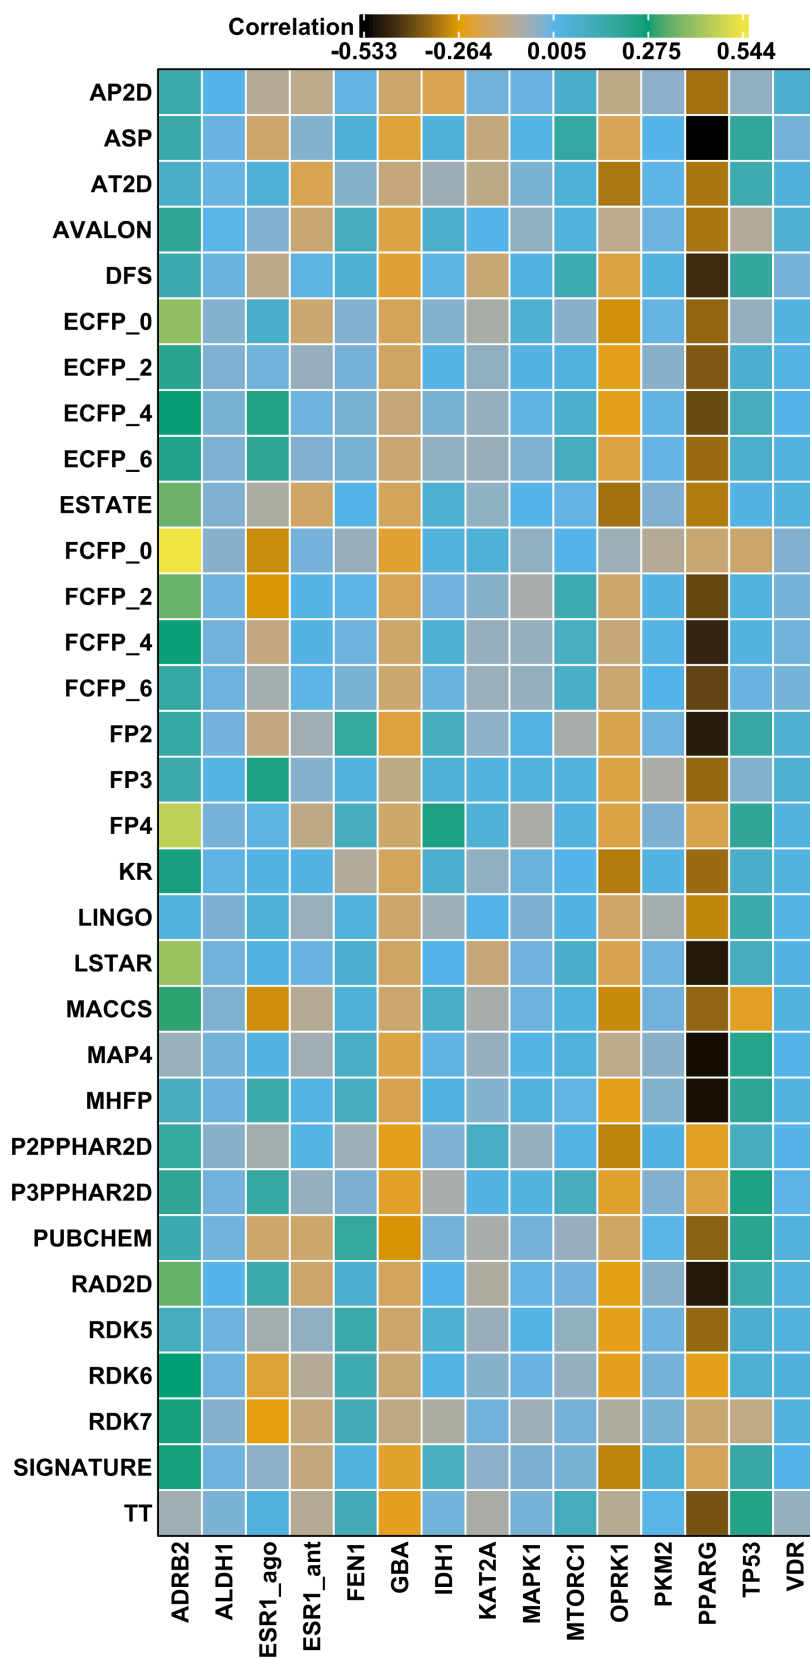


Figure 5. Heatmap of the Kendall rank correlation (τ) between fingerprint Tanimoto (T_c) similarities calculated between the most active compound for a given target and the potency values (AC_{50}) of the actives for that target.




Figure 6. Scatter plot of the fingerprint similarities (T_c) and the potencies (AC_{50}) of active compounds for ADRB2 (using FCFP0 fingerprint), VDR (SIGNATURE fingerprint), and PPARG (ASP fingerprint).

3. Discussion

Computational methods hold the promise of expanding the diversity of small molecule drug candidate targeting for specified protein binding pockets. To fulfill this promise, these methods must be supported by succinct molecular representation schemes that enable the rapid identification of molecules that are functionally similar to known or suspected active molecules while presenting non-trivial molecular similarity. The results of this study demonstrate that molecular fingerprints, and specifically the measurement of molecular similarity based on those fingerprints, are not generally effective at discriminating molecules with similar binding activity from those with dissimilar activity; furthermore, we have shown that where some early enrichment is observed, it tends to highlight candidate molecules that would arise from a simple structural perturbation. In addition to explicitly highlighting the weakness of molecular fingerprints for drug similarity assessment, the results also suggest that fingerprints are unlikely to be useful for the calculation of classification confidence or applicability domain [51]. In total, these results suggest that the task of the rapid evaluation of molecular similarity must move beyond the fingerprint representation of molecules.

The binding of a ligand to a target is a function of the target protein, the ligand, and other nearby molecules. In complicated cases, the role of nearby molecules can include cases where cooperation between multiple molecules is required for a binding event to occur. The ensemble properties of all the participant molecules should be considered to define similarity in the context of biological activity related to protein–ligand binding—a task difficult to achieve by focusing on ligand structure alone. The difficulty is poignantly highlighted by the fact that a simple atom count-based descriptor performs slightly worse than some fingerprint-based descriptors in defining similarity in biological activity [52]. When multiple active ligands are present for a target, with diverse chemical and structural properties differentiating consequential features from inconsequential ones in a binding event, fingerprint-based methods can possibly identify a novel active compound for that target. Multiple recent studies have shown success following that approach [53–55]. However, in most practical cases, such information is not available. It is important to understand

the limitation of defining “similarity” based on fingerprints in the context of biological activity and use such methods in conjunction with other orthogonal methods for a successful design of VS.

In recent years, deep learning strategies have entered the drug discovery toolkit (e.g., [56–58]), but these have not yet solved the problem of rapid virtual screening. Though the path forward is not clear, we suggest that it is vital that molecules be represented in such a way that the potential *context* of the molecule (i.e., information about the potential binding target) can be considered when evaluating the similarity of molecules. We suspect that future successful strategies will emphasize the surface properties of the small molecule [59,60], and will represent both the compound and the target protein not as a monolith, but as a collection of surface patches [61–63]. These, we believe, will enable a more context-dependent emphasis on features of importance to particular interactions, without interference from unimportant features.

4. Materials and Methods

4.1. Fingerprint Representations

The palette of fingerprints evaluated in this study (Table 3), can be broadly classified into those based on (i) path, (ii) circular features, (iii) pharmacophore, and (iv) pre-defined generic substructures/keys [64]. Circular and path-based fingerprints are generated using an exhaustive enumeration of (linear/circular) fragments up to a given size/radius, which are then hashed into a fixed-length bit vector. The SIGNATURE descriptor [65] generates explicitly defined substructures, which are mapped to numerical identifiers (no hashing involved). The LINGO fingerprint [16] works directly with the SMILES strings (rather than the usual dependence on a molecular graph representation), by fragmenting the strings into overlapping substrings.

Table 3. Molecular fingerprints evaluated in this study. Abbreviations: Topological torsion (TT), Extended Connectivity Fingerprint (ECFP), Functional Class Fingerprint (FCFP), Atom Pair (AP2D), Atom Triplet (AT2D), All Star Paths (ASP), Depth First Search (DFS).

TYPE	FAMILY	DESCRIPTION	SIZE (bits)
AP2D [66]	SUBSTRUCTURE	Topological Atom Pairs	4096
ASP [67]	PATH	All-Shortest Path encoding	4096
AT2D [67]	SUBSTRUCTURE	Topological Atom Triplets	4096
AVALON	SUBSTRUCTURE	Enumerates paths and feature classes	1024
DFS [68]	PATH	All-path encodings	4096
ECFP_0 [69,70]	CIRCULAR	Extended-connectivity fingerprint of diameter 0	1024
ECFP_2 [69,70]	CIRCULAR	Extended-connectivity fingerprint of diameter 2	1024
ECFP_4 [69,70]	CIRCULAR	Extended-connectivity fingerprint of diameter 4	1024
ECFP_6 [69,70]	CIRCULAR	Extended-connectivity fingerprint of diameter 6	1024
ESTATE [70,71]	SUBSTRUCTURE	Fingerprint based on E-State fragments	79
FCFP_0 [69,70]	CIRCULAR	Feature-class fingerprint of diameter 0	1024
FCFP_2 [69,70]	CIRCULAR	Feature-class fingerprint of diameter 2	1024
FCFP_4 [69,70]	CIRCULAR	Feature-class fingerprint of diameter 4	1024
FCFP_6 [69,70]	CIRCULAR	Feature-class fingerprint of diameter 6	1024
FP2 [72]	PATH	Indexes linear fragments up to 7 atoms in length	–
FP3 [72]	SUBSTRUCTURE	Based on 55 SMARTS patterns defining functional groups	–
FP4 [72]	SUBSTRUCTURE	Based on SMARTS patterns defining functional groups	–
KR [70,73]	SUBSTRUCTURE	Klekota–Roth SMARTS-based fingerprint	4860
LINGO [16,70]	TEXT	Fragmentation of SMILES strings	–
LSTAR [67]	PATH	Local Path Environments	4096
MACCS [74]	SUBSTRUCTURE	Molecular ACCess System structural keys	166
MAP4 [75]	CIRCULAR	Combines substructure and atom-pair concepts	2048
MHFP [76]	CIRCULAR	Encodes circular substructures	2048
P2PPHAR2D [77]	PHARMACOPHORE	Pharmacophore pair encoding	4096
P3PPHAR2D [77]	PHARMACOPHORE	Pharmacophore triplet encoding	4096
PUBCHEM [70,78]	SUBSTRUCTURE	Substructure fingerprint	881
RAD2D [22]	CIRCULAR	Topological Molprint-like fingerprints	4096
RDK5 [79]	SUBSTRUCTURE	Encodes substructures at most 5 bonds long	1024
RDK6 [79]	SUBSTRUCTURE	Encodes substructures at most 6 bonds long	1024
RDK7 [79]	SUBSTRUCTURE	Encodes substructures at most 7 bonds long	1024
SIGNATURE [65,70]	SUBSTRUCTURE	Based on an array of atom signatures	–
TT [66]	PATH	Based on bond paths of four non-hydrogen atoms	–

All fingerprints were generated using open-source software. Routines in the RDKit [79] library were used to compute the AVALON, ERG, RDK5, RDK6, RDK7, MHFP, and TT fingerprints. FP2, FP3 and FP4 fingerprint similarities were calculated directly using the OpenBabel toolbox [72]. The other fingerprints were calculated using custom software that makes use of the *jCompoundMapper* [67] and Chemistry Development Kit [70] libraries.

Although a number of similarity metrics have been used [80], the most common measure of fingerprint similarity is the Tanimoto coefficient [81]: $T_{ab} = |F_a \cap F_b| / |F_a \cup F_b|$, where F_a and F_b are the fingerprints of molecules a and b , respectively—this is equivalent to the Jaccard index over fingerprint bit vectors. Tanimoto coefficient values range between 1 (identical fingerprints, though not necessarily identical compounds) and 0 (disjoint fingerprints).

4.2. Benchmarking Data Sets

Numerous benchmarking data sets have been developed over the years to evaluate VS methods [39,82]. Each data set contains a set of active compounds (with known/documented activity for the target of interest) and a corresponding set of inactives/decoys. While the definition of actives is consistent, there is some variance in the question of what should be considered a ‘decoy’. Some benchmarks include only confirmed inactive molecules, while others add compounds presumed to be non-binding [39,83,84]. Data set composition can impact VS evaluation, such that both the artificial under- and over-estimation of enrichment have been well documented [39,85,86]. To account for benchmark-specific biases and error profiles, we opted to explore fingerprint efficacy across the full set of benchmarks. This evaluation pool provides a diverse perspective on the performance of molecular fingerprints, indicating that limited ability to differentiate decoys from active molecules is not simply due to a specific design flaw found in a single benchmark.

In this study, we employ four different VS data sets to explore the utility of molecular fingerprinting strategies for prediction of similar activity. These data sets are briefly summarized in Table 4 and described below.

Table 4. Comparison between different VS data sets. In all cases, the actives may not bind to the same pocket of the target.

Data Set	Active Source	Decoy Generation	Comments
DUD-E	ChEMBL09	Total of 0.65% from experiments. Total of 99.35% generated choosing different topologies with similar chemical properties using 2D similarity methods.	No rigorous method to remove false positives. Decoys biased towards 2D similarity methods.
DEKOIS	DUD (from literature) [87]	Decoys generated choosing different topologies with similar chemical properties using 2D similarity methods.	Low active to decoy ratio. Decoys biased towards 2D similarity methods
MUV	PubChem BioAssay	Unbiased distribution of decoys from experimentally available data.	Data processed to remove false positives and assay artifacts [88]. Low active-to-decoy ratio.
LIT-PCBA	PubChem BioAssay	Decoys were chosen from experimentally available data and pruned to have chemical properties similar to actives.	High active to decoy ratio. Actives may not bind to the same pocket of a target. Variable performance in 2D and 3D similarity search and docking across different target sets. Data processed to remove false positives and assay artifacts [88].

DUD-E: Directory of Useful Decoys, Enhanced [89]: DUD-E is a widely used data set for VS benchmarking, containing data for 102 protein targets. Each target is represented by an average of ~224 active ligands, ~90 experimental decoys, and ~14,000 computational decoys. Compounds are considered active based on a 1 μM experimental affinity cutoff, and experimental decoys are ligands with no measurable affinity up to 30 μM . Computational decoy ligands are selected from ZINC [83] to have 50 physical properties (rotatable bonds, hydrogen bond acceptors/donors, molecular weight, logP, net charge) similar to the actives, but with low fingerprint (Daylight [90]) Tanimoto coefficient $T_{ab} < 0.5$. (Note: this means that computational decoys are, by construction, expected to have low fingerprint similarity).

MUV: Maximum Unbiased Validation [91]: MUV data sets are based on bioactivity data available in PubChem [48]. This benchmark consists of sets of 30 actives (taken from confirmation assays) and 15,000 decoys (drawn from corresponding primary screens) for each of the 17 targets. The goal of the experimental design is to obtain an optimal spread of actives in the chemical space of the decoys. Since the data are taken from high-throughput screening assays that can be affected by experimental noise and artifacts (caused by unspecific activity of chemical compounds), an assay filter is applied to remove compounds interfering with optical detection methods (autofluorescence and luciferase inhibition) and potential aggregators.

DEKOIS: The Demanding Evaluation Kits for Objective In silico Screening (DEKOIS) [42] benchmark is based on BindingDB [92] bioactivity data (K_i , K_d , or IC_{50} values). The DEKOIS data set is derived from a set of 15 million molecules randomly selected from ZINC, which are divided into 10,752 bins based on their molecular weight (12 bins), octanol–water partition coefficient (8 bins), number of hydrogen bond acceptors (4 bins), number of hydrogen bond donors (4 bins), and number of rotatable bonds (7 bins). Active ligands are also placed into these pre-defined bins. For each active ligand, 1500 decoys are sampled from the active's bin (or neighboring bins, if necessary). These are further refined to a final set of 30 structurally diverse decoys per active. The DEKOIS data set includes 81 protein targets found in the DUD-E data set.

LIT-PCBA: The LIT-PCBA benchmark [43] is a curated subset of the PubChem BioAssay database [48], containing data from experiments where more than 10,000 chemicals were screened against a single protein target, and dose–response curves identified at least 50 actives. Active ligands identified in a bioassay experiment are not guaranteed to bind to the same pocket of the target protein; to overcome this concern, LIT-PCBA includes only targets with representative ligand-bound structures present in the PDB, such that the PDB ligands share the same phenotype or function as the true active ligands from the bioassay experiments. The LIT-PCBA data set was further refined to contain only targets for which at least one of the VS methods (2D fingerprint similarity, 3D shape similarity, and molecular docking) achieved an enrichment in true positives (i.e., the most challenging protein targets have been removed, so that enrichment results are, by design, expected to show good results). Targets in the LIT-PCBA have a variable active to decoy ratio that ranges from as low as 1:20 to 1:19,000.

4.3. Virtual Screening Evaluation

A common measure for the efficacy of a method's discriminatory power depends on the receiver operating characteristic (ROC) curve, which plots the sensitivity of a method as a function of false labels [93]. If a classification method assigns better scores to all true matches (actives) than to any false matches (decoys), then the area under that curve (AUC) will be 1. A random classifier will have an AUC of 0.5.

AUC provides a measure of the sensitivity/specificity trade-off across the full sensitivity range, but medicinal chemists are typically more interested in the early recognition of active molecules [32,94], since there is little actionable value to gain in the recall of true actives buried among thousands (or more) of decoys. As an example, consider an imaginary method that assigns the highest scores to 10% of all active molecules, then afterwards loses

discriminative power and assigns essentially random scores to all remaining molecules (actives and decoys). The AUC for such a method would be not particularly good (roughly 0.6), even though the early enrichment (ranking 10% of actives with a superior score to all decoys) provides some experimental utility.

To address this shortcoming of ROC AUC, a number of other metrics have been devised to assess early enrichment [32,95,96]. Unfortunately, it can be difficult to extract an intuitive meaning from these measures [97], and they are often not comparable across test sets because their scale and value depends on the set size and number of decoys in the test set. Here, we introduce a simple new early enrichment measure, the *decoy retention factor* (DRF); DRF is easy to interpret, and generalizes across input size. We note that DRF is only applicable in situations in which the number of active and decoy ligands is known beforehand. For the analysis of fingerprint benchmarks, we present both DRF and AUC values. Additional metrics such as BEDROC [94] and sum of log rank [32] are summarized visually in the Supplementary Information.

The purpose of DRF is to identify, for a parameterized fraction p of the active molecules, how effectively decoys are filtered from the score range containing those actives. Consider an input containing n active compounds and d decoys, and an enrichment threshold of $p = 0.1$. Since we are interested in the score of the top p fraction of actives, let $x = \lceil pn \rceil$, and let s_p be the score of the x -th element (so that s_p is the score threshold that would recover 10% of actives). Define d_p to be the number of decoys that exceed s_p —this is a fraction of the d total decoys. DRF_p measures the extent to which decoys have been filtered out of the range containing the top p actives:

$$DRF_p = \frac{d_p}{\lceil pd \rceil} \quad (1)$$

A system that assigns scores randomly will recover a fraction p of the decoys at roughly the same score as it sees p of the actives, so that $DRF_p = 1$. In an ideal case, no decoys have score greater than the x -th active element, meaning that $d_p = 0$, and thus $DRF_p = 0$. A $DRF_p = 0.2$ indicates that the number of decoys remaining is 20% of the number expected if p of the decoys were kept (there is a 5-fold reduction in decoys). Meanwhile a $DRF_p > 1$ indicates that the method *enriches* for decoys.

We find DRF to be a useful measure because it enables the prediction of the number of decoys expected to remain in a score-filtered result set, based on the size of the underlying library. For example, consider a library of 1 million molecules—this will consist almost entirely of inactives (decoys), so that $d \approx 1,000,000$. If we hope to discover 10% of actives, and we have previously established that $DRF_{0.1} = 0.05$ (a 20-fold reduction in decoys relative to random chance), then we expect to observe $d_p \cdot DRF_p \approx 1,000,000 \cdot 0.1 \cdot 0.05 = 5000$ decoys mixed with the surviving actives. This simple calculation is important, because it can highlight that apparently good enrichment (10- or 100-fold) may not be enough to effectively filter out inactives when the target set includes billions of candidates.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/ph17080992/s1>, Table S1: Summary of the VS screening performances using extended length fingerprints; Figure S1: Heatmap of the area under the curve (AUC) obtained by the different fingerprints for the targets in the benchmark datasets; Figure S2: Heatmap of the decoy retention factor ($DRF_{0.1}$) obtained by the different fingerprints for the targets in the benchmark datasets; Figure S3: Heatmap of the BEDROC ($\alpha = 20$) values obtained by the different fingerprints for the targets in the benchmark datasets; Figure S4: Heatmap of the enrichment factor (top 8%) obtained by the different fingerprints for the targets in the benchmark datasets; Figure S5: Heatmap of the NSLR values obtained by the different fingerprints for the targets in the benchmark datasets; Figure S6: In the LIT-PCBA analysis, early enrichment was observed for PPARG; Figure S7: In the LIT-PCBA analysis, early enrichment was observed for GBA. [32,94,97]

Author Contributions: Conceptualization, V.V., A.R. and T.J.W.; methodology, V.V., A.R. and T.J.W.; software, V.V., J.G., D.D., A.R. and T.J.W.; validation, V.V.; investigation, V.V., D.D. and R.X.; resources, T.J.W.; data curation, V.V.; writing—original draft preparation, V.V. and T.J.W.; writing—review and editing, all; visualization, V.V., D.D. and R.X.; supervision, T.J.W.; project administration, T.J.W.; funding acquisition, T.J.W. All authors have read and agreed to the published version of the manuscript.

Funding: VV was supported by the Research Council of Norway through grant no. 275752. AR was supported by the NIH National Institute of Allergy and Infectious Diseases (NIAID), Department of Health and Human Services under BCBB Support Services Contract HHSN316201300006W/HHSN27200002 to MSC, Inc. JLG was supported by the Federal Pathway Program at NIH, NIAID. RX was supported by NIH National Institute of Allergy and Infectious Diseases (NIAID) grant AI168165. TJW and DD were supported by NIH National Institute of General Medical Sciences (NIGMS) grant GM132600, and by the DOE Biological and Environmental Research (BER) Program, grant DE-SC0021216.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data and software used for the calculation of fingerprints and scripts to reproduce the results are available from <https://osf.io/d3cbr/>.

Acknowledgments: We gratefully acknowledge the computational resources and expert administration provided by the University of Montana's Griz Shared Computing Cluster (GSCC) and the high performance computing (HPC) resources supported by the University of Arizona TRIF, UITS, and Research, Innovation, and Impact (RII) and maintained by the UArizona Research Technologies department.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

VS	Virtual Screening
SPP	Similar Property Principle
LBVS	Ligand-Based Virtual Screening
DRF	Decoy Retention Factor
ROC	Receiver Operating Characteristic
AUC	Area Under the Curve

References

1. Sadybekov, A.A.; Sadybekov, A.V.; Liu, Y.; Iliopoulos-Tsoutsouvas, C.; Huang, X.P.; Pickett, J.; Houser, B.; Patel, N.; Tran, N.K.; Tong, F.; et al. Synthon-based ligand discovery in virtual libraries of over 11 billion compounds. *Nature* **2022**, *601*, 452–459. [[CrossRef](#)]
2. Venkatraman, V.; Colligan, T.H.; Lesica, G.T.; Olson, D.R.; Gaiser, J.; Copeland, C.J.; Wheeler, T.J.; Roy, A. Drugsniffer: An Open Source Workflow for Virtually Screening Billions of Molecules for Binding Affinity to Protein Targets. *Front. Pharmacol.* **2022**, *13*, 874746. [[CrossRef](#)]
3. Luttens, A.; Gullberg, H.; Abdurakhmanov, E.; Vo, D.D.; Akaberi, D.; Talibov, V.O.; Nekhotiaeva, N.; Vangeel, L.; De Jonghe, S.; Jochmans, D.; et al. Ultralarge virtual screening identifies SARS-CoV-2 main protease inhibitors with broad-spectrum activity against coronaviruses. *J. Am. Chem. Soc.* **2022**, *144*, 2905–2920. [[CrossRef](#)]
4. Warr, W.A.; Nicklaus, M.C.; Nicolaou, C.A.; Rarey, M. Exploration of ultralarge compound collections for drug discovery. *J. Chem. Inf. Model.* **2022**, *62*, 2021–2034. [[CrossRef](#)]
5. Walters, W.P. Virtual Chemical Libraries. *J. Med. Chem.* **2018**, *62*, 1116–1124. [[CrossRef](#)]
6. Gimeno, A.; Ojeda-Montes, M.; Tomás-Hernández, S.; Cereto-Massagué, A.; Beltrán-Debón, R.; Mulero, M.; Pujadas, G.; Garcia-Vallvé, S. The Light and Dark Sides of Virtual Screening: What Is There to Know? *Int. J. Mol. Sci.* **2019**, *20*, 1375. [[CrossRef](#)]
7. Maia, E.H.B.; Assis, L.C.; de Oliveira, T.A.; da Silva, A.M.; Taranto, A.G. Structure-Based Virtual Screening: From Classical to Artificial Intelligence. *Front. Chem.* **2020**, *8*, 343. [[CrossRef](#)]
8. Johnson, M.; Maggiora, G.M. *Concepts and Applications of Molecular Similarity*; Wiley: New York, NY, USA, 1990.
9. Maggiora, G.; Vogt, M.; Stumpfe, D.; Bajorath, J. Molecular Similarity in Medicinal Chemistry. *J. Med. Chem.* **2013**, *57*, 3186–3204. [[CrossRef](#)]
10. Cortés-Ciriano, I.; Škuta, C.; Bender, A.; Svozil, D. QSAR-derived affinity fingerprints (part 2): Modeling performance for potency prediction. *J. Cheminf.* **2020**, *12*, 41. [[CrossRef](#)]

11. Venkatraman, V. FP-MAP: An extensive library of fingerprint-based molecular activity prediction tools. *Front. Chem.* **2023**, *11*, 1239467. [[CrossRef](#)]
12. Tetko, I.V.; Bruneau, P.; Mewes, H.W.; Rohrer, D.C.; Poda, G.I. Can we estimate the accuracy of ADME–Tox predictions? *Drug Discov. Today* **2006**, *11*, 700–707. [[CrossRef](#)]
13. Mellor, C.; Robinson, R.M.; Benigni, R.; Ebbrell, D.; Enoch, S.; Firman, J.; Madden, J.; Pawar, G.; Yang, C.; Cronin, M. Molecular fingerprint-derived similarity measures for toxicological read-across: Recommendations for optimal use. *Regul. Toxicol. Pharmacol.* **2019**, *101*, 121–134. [[CrossRef](#)]
14. Venkatraman, V. FP-ADMET: A compendium of fingerprint-based ADMET prediction models. *J. Cheminf.* **2021**, *13*, 75. [[CrossRef](#)]
15. Teixeira, A.L.; Falcao, A.O. Structural Similarity Based Kriging for Quantitative Structure Activity and Property Relationship Modeling. *J. Chem. Inf. Model.* **2014**, *54*, 1833–1849. [[CrossRef](#)]
16. Vidal, D.; Thormann, M.; Pons, M. LINGO, an Efficient Holographic Text Based Method To Calculate Biophysical Properties and Intermolecular Similarities. *J. Chem. Inf. Model.* **2005**, *45*, 386–393. [[CrossRef](#)]
17. Kubinyi, H. Similarity and dissimilarity: A medicinal chemist’s view. *Perspect. Drug Discov. Des.* **1998**, *9*, 225–252. [[CrossRef](#)]
18. Maggiora, G.M. On Outliers and Activity Cliffs Why QSAR Often Disappoints. *J. Chem. Inf. Model.* **2006**, *46*, 1535–1535. [[CrossRef](#)]
19. Stumpfe, D.; Hu, H.; Bajorath, J. Evolving Concept of Activity Cliffs. *ACS Omega* **2019**, *4*, 14360–14368. [[CrossRef](#)]
20. Van Tilborg, D.; Alenicheva, A.; Grisoni, F. Exposing the limitations of molecular machine learning with activity cliffs. *J. Chem. Inf. Model.* **2022**, *62*, 5938–5951. [[CrossRef](#)]
21. Barbosa, F.; Horvath, D. Molecular Similarity and Property Similarity. *Curr. Top. Med. Chem.* **2004**, *4*, 589–600. [[CrossRef](#)]
22. Bender, A.; Glen, R.C. Molecular similarity: A key technique in molecular informatics. *Org. Biomol. Chem.* **2004**, *2*, 3204. [[CrossRef](#)] [[PubMed](#)]
23. Muratov, E.N.; Bajorath, J.; Sheridan, R.P.; Tetko, I.V.; Filimonov, D.; Poroikov, V.; Oprea, T.I.; Baskin, I.I.; Varnek, A.; Roitberg, A.; et al. QSAR without borders. *Chem. Soc. Rev.* **2020**, *49*, 3525–3564. [[CrossRef](#)] [[PubMed](#)]
24. Willett, P. Similarity Searching Using 2D Structural Fingerprints. In *Methods in Molecular Biology*; Humana Press: Totowa, NJ, USA, 2010; pp. 133–158. [[CrossRef](#)]
25. Stumpfe, D.; Bajorath, J. Similarity searching. *WIREs Comput. Mol. Sci.* **2011**, *1*, 260–282. [[CrossRef](#)]
26. Cereto-Massagué, A.; Ojeda, M.J.; Valls, C.; Mulero, M.; Garcia-Vallvé, S.; Pujadas, G. Molecular fingerprint similarity search in virtual screening. *Methods* **2015**, *71*, 58–63. [[CrossRef](#)]
27. Gao, K.; Nguyen, D.D.; Sresht, V.; Mathiowetz, A.M.; Tu, M.; Wei, G.W. Are 2D fingerprints still valuable for drug discovery? *Phys. Chem. Chem. Phys.* **2020**, *22*, 8373–8390. [[CrossRef](#)] [[PubMed](#)]
28. Martin, Y.C.; Kofron, J.L.; Traphagen, L.M. Do Structurally Similar Molecules Have Similar Biological Activity? *J. Med. Chem.* **2002**, *45*, 4350–4358. [[CrossRef](#)] [[PubMed](#)]
29. Shoichet, B.K. Virtual screening of chemical libraries. *Nature* **2004**, *432*, 862–865. [[CrossRef](#)] [[PubMed](#)]
30. Hann, M.M.; Oprea, T.I. Pursuing the leadlikeness concept in pharmaceutical research. *Curr. Opin. Chem. Biol.* **2004**, *8*, 255–263. [[CrossRef](#)] [[PubMed](#)]
31. Leeson, P.D.; Springthorpe, B. The influence of drug-like concepts on decision-making in medicinal chemistry. *Nat. Rev. Drug Discov.* **2007**, *6*, 881–890. [[CrossRef](#)]
32. Venkatraman, V.; Pérez-Nuño, V.I.; Mavridis, L.; Ritchie, D.W. Comprehensive Comparison of Ligand-Based Virtual Screening Tools Against the DUD Data set Reveals Limitations of Current 3D Methods. *J. Chem. Inf. Model.* **2010**, *50*, 2079–2093. [[CrossRef](#)]
33. Sciabola, S.; Torella, R.; Nagata, A.; Boehm, M. Critical Assessment of State-of-the-Art Ligand-Based Virtual Screening Methods. *Mol. Inf.* **2022**, *41*, 2200103. [[CrossRef](#)] [[PubMed](#)]
34. Verdonk, M.L.; Berdini, V.; Hartshorn, M.J.; Mooij, W.T.M.; Murray, C.W.; Taylor, R.D.; Watson, P. Virtual Screening Using Protein-Ligand Docking: Avoiding Artificial Enrichment. *J. Chem. Inf. Model.* **2004**, *44*, 793–806. [[CrossRef](#)]
35. Sieg, J.; Flachsenberg, F.; Rarey, M. In Need of Bias Control: Evaluating Chemical Data for Machine Learning in Structure-Based Virtual Screening. *J. Chem. Inf. Model.* **2019**, *59*, 947–961. [[CrossRef](#)] [[PubMed](#)]
36. Li, H.; Sze, K.H.; Lu, G.; Ballester, P.J. Machine-learning scoring functions for structure-based drug lead optimization. *WIREs Comput. Mol. Sci.* **2020**, *10*, e1465. [[CrossRef](#)]
37. Imrie, F.; Bradley, A.R.; Deane, C.M. Generating property-matched decoy molecules using deep learning. *Bioinformatics* **2021**, *37*, 2134–2141. [[CrossRef](#)] [[PubMed](#)]
38. Stein, R.M.; Yang, Y.; Balius, T.E.; O’Meara, M.J.; Lyu, J.; Young, J.; Tang, K.; Shoichet, B.K.; Irwin, J.J. Property-Unmatched Decoys in Docking Benchmarks. *J. Chem. Inf. Model.* **2021**, *61*, 699–714. [[CrossRef](#)] [[PubMed](#)]
39. Réau, M.; Langenfeld, F.; Zagury, J.F.; Lagarde, N.; Montes, M. Decoys Selection in Benchmarking Datasets: Overview and Perspectives. *Front. Pharmacol.* **2018**, *9*, 11. [[CrossRef](#)] [[PubMed](#)]
40. Chen, L.; Cruz, A.; Ramsey, S.; Dickson, C.J.; Duca, J.S.; Hornak, V.; Koes, D.R.; Kurtzman, T. Hidden bias in the DUD-E dataset leads to misleading performance of deep learning in structure-based virtual screening. *PLoS ONE* **2019**, *14*, e0220113. [[CrossRef](#)] [[PubMed](#)]
41. O’Boyle, N.M.; Sayle, R.A. Comparing structural fingerprints using a literature-based similarity benchmark. *J. Cheminf.* **2016**, *8*, 36. [[CrossRef](#)]
42. Bauer, M.R.; Ibrahim, T.M.; Vogel, S.M.; Boeckler, F.M. Evaluation and Optimization of Virtual Screening Workflows with DEKOIS 2.0—A Public Library of Challenging Docking Benchmark Sets. *J. Chem. Inf. Model.* **2013**, *53*, 1447–1462. [[CrossRef](#)]

43. Tran-Nguyen, V.K.; Jacquemard, C.; Rognan, D. LIT-PCBA: An Unbiased Data Set for Machine Learning and Virtual Screening. *J. Chem. Inf. Model.* **2020**, *60*, 4263–4273. [[CrossRef](#)]
44. Baldi, P.; Nasr, R. When is chemical similarity significant? The statistical distribution of chemical similarity scores and its extreme values. *J. Chem. Inf. Model.* **2010**, *50*, 1205–1222. [[CrossRef](#)]
45. O Hagan, S.; Swainston, N.; Handl, J.; Kell, D.B. A rule of 0.5 for the metabolite-likeness of approved pharmaceutical drugs. *Metabolomics* **2015**, *11*, 323–339. [[CrossRef](#)]
46. Vogt, M.; Bajorath, J. Modeling tanimoto similarity value distributions and predicting search results. *Mol. Inform.* **2017**, *36*, 1600131. [[CrossRef](#)]
47. Verras, A.; Waller, C.L.; Gedeck, P.; Green, D.V.S.; Kogej, T.; Raichurkar, A.; Panda, M.; Shelat, A.A.; Clark, J.; Guy, R.K.; et al. Shared Consensus Machine Learning Models for Predicting Blood Stage Malaria Inhibition. *J. Chem. Inf. Model.* **2017**, *57*, 445–453. [[CrossRef](#)]
48. Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B.A.; Thiessen, P.A.; Yu, B.; et al. PubChem 2023 update. *Nucleic Acids Res.* **2022**, *51*, D1373–D1380. [[CrossRef](#)]
49. Kendall, M.G. A new measure of rank correlation. *Biometrika* **1938**, *30*, 81–93. [[CrossRef](#)]
50. Vogt, M.; Bajorath, J. Similarity Searching for Potent Compounds Using Feature Selection. *J. Chem. Inf. Model.* **2013**, *53*, 1613–1619. [[CrossRef](#)]
51. Mathea, M.; Klingspohn, W.; Baumann, K. Chemoinformatic classification methods and their applicability domain. *Mol. Inform.* **2016**, *35*, 160–180. [[CrossRef](#)]
52. Bender, A. How similar are those molecules after all? Use two descriptors and you will have three different answers. *Expert Opin. Drug Discov.* **2010**, *5*, 1141–1151. [[CrossRef](#)]
53. Kilchmann, F.; Marcaida, M.J.; Kotak, S.; Schick, T.; Boss, S.D.; Awale, M.; Gonczy, P.; Reymond, J.L. Discovery of a selective aurora a kinase inhibitor by virtual screening. *J. Med. Chem.* **2016**, *59*, 7188–7211. [[CrossRef](#)]
54. Ozhathil, L.C.; Delalande, C.; Bianchi, B.; Nemeth, G.; Kappel, S.; Thomet, U.; Ross-Kaschitza, D.; Simonin, C.; Rubin, M.; Gertsch, J.; et al. Identification of potent and selective small molecule inhibitors of the cation channel TRPM4. *Br. J. Pharmacol.* **2018**, *175*, 2504–2519. [[CrossRef](#)]
55. Zhu, D.; Johannsen, S.; Masini, T.; Simonin, C.; Hauptenthal, J.; Illarionov, B.; Andreas, A.; Awale, M.; Gierse, R.M.; van der Laan, T.; et al. Discovery of novel drug-like antitubercular hits targeting the MEP pathway enzyme DXPS by strategic application of ligand-based virtual screening. *Chem. Sci.* **2022**, *13*, 10686–10698. [[CrossRef](#)]
56. Yang, Y.; Yao, K.; Repasky, M.P.; Leswing, K.; Abel, R.; Shoichet, B.K.; Jerome, S.V. Efficient exploration of chemical space with docking and deep learning. *J. Chem. Theory Comput.* **2021**, *17*, 7106–7119. [[CrossRef](#)]
57. Gorantla, R.; Kubincova, A.; Weiße, A.Y.; Mey, A.S. From Proteins to Ligands: Decoding Deep Learning Methods for Binding Affinity Prediction. *J. Chem. Inf. Model.* **2023**, *64*, 2496–2507. [[CrossRef](#)]
58. Gentile, F.; Yaacoub, J.C.; Gleave, J.; Fernandez, M.; Ton, A.T.; Ban, F.; Stern, A.; Cherkasov, A. Artificial intelligence-enabled virtual screening of ultra-large chemical libraries with deep docking. *Nat. Protoc.* **2022**, *17*, 672–697. [[CrossRef](#)]
59. Pérez-Nueno, V.I.; Venkatraman, V.; Mavridis, L.; Clark, T.; Ritchie, D.W. Using Spherical Harmonic Surface Property Representations for Ligand-Based Virtual Screening. *Mol. Inform.* **2010**, *30*, 151–159. [[CrossRef](#)]
60. Pérez-Nueno, V.I.; Venkatraman, V.; Mavridis, L.; Ritchie, D.W. Detecting Drug Promiscuity Using Gaussian Ensemble Screening. *J. Chem. Inf. Model.* **2012**, *52*, 1948–1961. [[CrossRef](#)]
61. Hofbauer, C.; Lohninger, H.; Aszódi, A. SURFCOMP: A Novel Graph-Based Approach to Molecular Surface Comparison. *J. Chem. Inf. Model.* **2004**, *44*, 837–847. [[CrossRef](#)]
62. Gainza, P.; Sverrisson, F.; Monti, F.; Rodolà, E.; Boscaini, D.; Bronstein, M.M.; Correia, B.E. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nat. Methods* **2019**, *17*, 184–192. [[CrossRef](#)]
63. Douguet, D.; Payan, F. sensaas: Shape-based Alignment by Registration of Colored Point-based Surfaces. *Mol. Inf.* **2020**, *39*, 2000081. [[CrossRef](#)] [[PubMed](#)]
64. Bender, A.; Jenkins, J.L.; Scheiber, J.; Sukuru, S.C.K.; Glick, M.; Davies, J.W. How Similar Are Similarity Searching Methods? A Principal Component Analysis of Molecular Descriptor Space. *J. Chem. Inf. Model.* **2009**, *49*, 108–119. [[CrossRef](#)] [[PubMed](#)]
65. Faulon, J.L.; Visco, D.P.; Pophale, R.S. The Signature Molecular Descriptor. 1. Using Extended Valence Sequences in QSAR and QSPR Studies. *J. Chem. Inf. Model.* **2003**, *43*, 707–720. [[CrossRef](#)] [[PubMed](#)]
66. Carhart, R.E.; Smith, D.H.; Venkataraghavan, R. Atom pairs as molecular features in structure-activity studies: Definition and applications. *J. Chem. Inf. Model.* **1985**, *25*, 64–73. [[CrossRef](#)]
67. Hinselmann, G.; Rosenbaum, L.; Jahn, A.; Fechner, N.; Zell, A. jCompoundMapper: An open source Java library and command-line tool for chemical fingerprints. *J. Cheminf.* **2011**, *3*, 3. [[CrossRef](#)]
68. Ralaivola, L.; Swamidass, S.J.; Saigo, H.; Baldi, P. Graph kernels for chemical informatics. *Neural Netw.* **2005**, *18*, 1093–1110. [[CrossRef](#)] [[PubMed](#)]
69. Rogers, D.; Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754. [[CrossRef](#)]
70. Willighagen, E.L.; Mayfield, J.W.; Alvarsson, J.; Berg, A.; Carlsson, L.; Jeliaskova, N.; Kuhn, S.; Pluskal, T.; Rojas-Chertó, M.; Spjuth, O.; et al. The Chemistry Development Kit (CDK) v2.0: Atom typing, depiction, molecular formulas, and substructure searching. *J. Cheminf.* **2017**, *9*, 33. [[CrossRef](#)] [[PubMed](#)]

71. Hall, L.H.; Kier, L.B. Electrotopological State Indices for Atom Types: A Novel Combination of Electronic, Topological, and Valence State Information. *J. Chem. Inf. Model.* **1995**, *35*, 1039–1045. [[CrossRef](#)]
72. O'Boyle, N.M.; Banck, M.; James, C.A.; Morley, C.; Vandermeersch, T.; Hutchison, G.R. Open Babel: An open chemical toolbox. *J. Cheminf.* **2011**, *3*, 33. [[CrossRef](#)]
73. Klekota, J.; Roth, F.P. Chemical substructures that enrich for biological activity. *Bioinformatics* **2008**, *24*, 2518–2525. [[CrossRef](#)] [[PubMed](#)]
74. Durant, J.L.; Leland, B.A.; Henry, D.R.; Nourse, J.G. Reoptimization of MDL Keys for Use in Drug Discovery. *J. Chem. Inf. Model.* **2002**, *42*, 1273–1280. [[CrossRef](#)] [[PubMed](#)]
75. Capecchi, A.; Probst, D.; Reymond, J.L. One molecular fingerprint to rule them all: Drugs, biomolecules, and the metabolome. *J. Cheminf.* **2020**, *12*, 43. [[CrossRef](#)] [[PubMed](#)]
76. Probst, D.; Reymond, J.L. A probabilistic molecular fingerprint for big data settings. *J. Cheminf.* **2018**, *10*, 66. [[CrossRef](#)] [[PubMed](#)]
77. Mahé, P.; Ralaivola, L.; Stoven, V.; Vert, J.P. The Pharmacophore Kernel for Virtual Screening with Support Vector Machines. *J. Chem. Inf. Model.* **2006**, *46*, 2003–2014. [[CrossRef](#)] [[PubMed](#)]
78. PubChem Substructure Fingerprint, Version: 1.3. 2022. Available online: <ftp://ftp.ncbi.nlm.nih.gov/pubchem/specifications/> (accessed on 1 March 2024).
79. Landrum, G. RDKit: Open-Source Cheminformatics, 2022. Release: 2022.03.5. Available online: <https://www.rdkit.org> (accessed on 1 March 2024).
80. Raymond, J.W.; Willett, P. Effectiveness of graph-based and fingerprint-based similarity measures for virtual screening of 2D chemical structure databases. *J. Comput. Aided Mol. Des.* **2002**, *16*, 59–71. [[CrossRef](#)] [[PubMed](#)]
81. Bajusz, D.; Rácz, A.; Héberger, K. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *J. Cheminf.* **2015**, *7*, 20. [[CrossRef](#)] [[PubMed](#)]
82. Lagarde, N.; Zagury, J.F.; Montes, M. Benchmarking Data Sets for the Evaluation of Virtual Ligand Screening Methods: Review and Perspectives. *J. Chem. Inf. Model.* **2015**, *55*, 1297–1307. [[CrossRef](#)]
83. Irwin, J.J. Community benchmarks for virtual screening. *J. Comput. Aided Mol. Des.* **2008**, *22*, 193–199. [[CrossRef](#)]
84. Tran-Nguyen, V.K.; Rognan, D. Benchmarking Data Sets from PubChem BioAssay Data: Current Scenario and Room for Improvement. *Int. J. Mol. Sci.* **2020**, *21*, 4380. [[CrossRef](#)]
85. Nisius, B.; Bajorath, J. Rendering Conventional Molecular Fingerprints for Virtual Screening Independent of Molecular Complexity and Size Effects. *ChemMedChem* **2010**, *5*, 859–868. [[CrossRef](#)] [[PubMed](#)]
86. Chaput, L.; Martinez-Sanz, J.; Saettel, N.; Mouawad, L. Benchmark of four popular virtual screening programs: Construction of the active/decoy dataset remains a major determinant of measured performance. *J. Cheminf.* **2016**, *8*, 56. [[CrossRef](#)] [[PubMed](#)]
87. Huang, N.; Shoichet, B.K.; Irwin, J.J. Benchmarking sets for molecular docking. *J. Med. Chem.* **2006**, *49*, 6789–6801. [[CrossRef](#)] [[PubMed](#)]
88. Feng, B.Y.; Shelat, A.; Doman, T.N.; Guy, R.K.; Shoichet, B.K. High-throughput assays for promiscuous inhibitors. *Nat. Chem. Biol.* **2005**, *1*, 146–148. [[CrossRef](#)]
89. Mysinger, M.M.; Carchia, M.; Irwin, J.J.; Shoichet, B.K. Directory of Useful Decoys, Enhanced (DUD-E): Better Ligands and Decoys for Better Benchmarking. *J. Med. Chem.* **2012**, *55*, 6582–6594. [[CrossRef](#)] [[PubMed](#)]
90. Daylight Chemical Information Systems, I. Daylight Theory Manual, Version 4.9. 2011. Available online: <https://www.daylight.com/dayhtml/doc/theory> (accessed on 1 March 2024).
91. Rohrer, S.G.; Baumann, K. Maximum Unbiased Validation (MUV) Data Sets for Virtual Screening Based on PubChem Bioactivity Data. *J. Chem. Inf. Model.* **2009**, *49*, 169–184. [[CrossRef](#)] [[PubMed](#)]
92. Gilson, M.K.; Liu, T.; Baitaluk, M.; Nicola, G.; Hwang, L.; Chong, J. BindingDB in 2015: A public database for medicinal chemistry, computational chemistry and systems pharmacology. *Nucleic Acids Res.* **2015**, *44*, D1045–D1053. [[CrossRef](#)] [[PubMed](#)]
93. Hanley, J.A.; McNeil, B.J. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* **1982**, *143*, 29–36. [[CrossRef](#)]
94. Truchon, J.F.; Bayly, C.I. Evaluating Virtual Screening Methods: Good and Bad Metrics for the “Early Recognition” Problem. *J. Chem. Inf. Model.* **2007**, *47*, 488–508. [[CrossRef](#)]
95. Clark, R.D.; Webster-Clark, D.J. Managing bias in ROC curves. *J. Comput. Aided Mol. Des.* **2008**, *22*, 141–146. [[CrossRef](#)]
96. Lopes, J.C.D.; dos Santos, F.M.; Martins-José, A.; Augustyns, K.; Winter, H.D. The power metric: A new statistically robust enrichment-type metric for virtual screening applications with early recovery capability. *J. Cheminf.* **2017**, *9*, 7. [[CrossRef](#)] [[PubMed](#)]
97. Zhao, W.; Hevener, K.E.; White, S.W.; Lee, R.E.; Boyett, J.M. A statistical framework to evaluate virtual screening. *BMC Bioinf.* **2009**, *10*, 225. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.