

001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046

Topological Regression as an interpretable and efficient tool for Quantitative Structure-Activity Relationship Modeling

Ruibo Zhang^{1†}, Daniel Nolte^{1†}, Cesar Sanchez-Villalobos¹,
Souparno Ghosh^{2*}, Ranadip Pal^{1*}

¹Dept. of Electrical and Computer Engineering, Texas Tech University,
Lubbock, TX 79409, USA.

²Department of Statistics, University of Nebraska - Lincoln, Lincoln,
NB 68588, USA.

*Corresponding author(s). E-mail(s): sghosh5@unl.edu;
ranadip.pal@ttu.edu;

Contributing authors: ruibo.zhang@ttu.edu; daniel.nolte@ttu.edu;
cesarasa@ttu.edu;

†These authors contributed equally.

Abstract

Quantitative structure-activity relationship (QSAR) modeling is a powerful tool for drug discovery, yet the lack of interpretability of commonly used QSAR models hinders their application in molecular design. We propose a similarity-based regression framework, topological regression (TR), that offers a statistically grounded, computationally fast, and interpretable technique to predict drug responses. We compare the predictive performance of TR on 530 ChEMBL human target activity datasets against the predictive performance of deep-learning-based QSAR models. Our results suggest that our sparse TR model can achieve equal, if not better, performance than the deep learning-based QSAR models and provide better intuitive interpretation by extracting an approximate isometry between the chemical space of the drugs and their activity space.

Keywords: Machine learning, Drug discovery, QSAR, Topological regression

047 1 Introduction

048

049 Quantitative structure-activity relationship (QSAR) models have become an essen-
050 tial tool in pharmaceutical discovery, especially in the virtual screening for hits and
051 lead optimization stages [1]. Experimental characterization of candidate molecules is
052 expensive and time-consuming. As a relatively easy-to-implement alternative, QSAR
053 models could be a valuable tool for assisting chemists by providing design ideas to
054 prioritize their experiments. QSARs are usually supervised machine learning models
055 that describe the connections between chemical structures and their biological activi-
056 ties, such as their potency, physicochemical properties, pharmacokinetic properties, or
057 environmental effects [2]. QSAR models enable *in silico* structural design by providing
058 property predictions from machine-readable representations of the chemical structure,
059 thereby helping generate and prioritize design ideas. This technique has been widely
060 applied in virtual screening and lead optimization with a fair amount of success [1, 3].

061 In QSAR methods, chemical substances must first be transformed into machine-
062 comprehensible mathematical representations. Three commonly used representations
063 are: (a) vectors such as classical molecular descriptors or molecular fingerprints (FPs),
064 (b) graphs, and (c) strings such as Simplified Molecular Input Line Entry System
065 (SMILES). Classical molecular descriptors[4] encode a specific computed or measured
066 attribute of the molecule into a single number, for instance, the count of bonds, atoms,
067 functional groups, or physicochemical characteristics, and are often used in combi-
068 nation to form feature vectors. PaDEL [5], Mordred [6], and RDKit are examples
069 of popular descriptor-calculation software packages for numerically representing the
070 chemical structure and molecular characteristics. Extended-connectivity fingerprints
071 (ECFPs)[7] are an example of a topological fingerprint computed using a variant of
072 the Morgan Algorithm that encodes chemical substructures by atom neighborhoods
073 using a high-dimensional sparse bit-string representation. The graph representation,
074 on the other hand, characterizes 2D chemical structures as graphs, with atoms as ver-
075 tices and bonds as edges. SMILES specify a notation for representing the chemical
076 graphs of molecules as strings of characters.

077 Once the chemical structures are represented using a suitable protocol, a predictive
078 method is chosen to connect the structural information with the functional properties.
079 For instance, if the chemical structures are represented as strings or graphs, deep-
080 learning methods are often used for prediction due to their ability to perform embedded
081 feature extraction. Chemprop [8], in particular, has turned out to be a popular method
082 that uses directed message-passing neural networks to learn molecular representations
083 directly from the graphs to predict the properties of molecules. This method has been
084 shown to excel at antibiotic discovery [9, 10] and lipophilicity prediction [11] indicating
085 its potential as a QSAR model. With the rise in popularity of large language models
086 and the attention mechanism, the use of SMILES strings has been increasingly investi-
087 gated for their potential embedded feature extraction, predictive performance, and
088 interpretability. For example, [12] pre-trained a transformer-based network through
089 masked SMILES recovery, and offered the pre-trained model for transfer learning
090 onto specific tasks. Similarly, Transformer-Convolutional Neural Network (CNN) [13]
091 applied the transformer architecture to canonicalize SMILES string inputs and enables
092 transfer learning of the model onto specific activity prediction tasks.

QSAR models are often developed for their predictive performance. However, the effectiveness of QSAR models, as a computational tool assisting molecular discovery and design, could be greatly improved by enhancing their domain-specific interpretability. Model interpretability, usually defined as the ability to explain predictions in a human-understandable way[14], typically consists of computing feature importance scores[15–18], influence functions to identify training instances most responsible for the prediction[19], developing locally interpretable models to approximate global black-box algorithms[20–22], and generating counterfactuals[23, 24]. For example, standard shallow learners, like Random Forests (RF) and Support Vector Machines (SVM) are often used in QSAR modeling to offer feature importance scores[25]. However, molecular interpretability is largely based on the interpretability of the underlying molecular representation. For instance, ALogP can be used as an important classical descriptor that plays a key role in determining the solubility of a molecule. However, a target value of ALogP cannot be mapped back to a precise chemical structure. When using interpretable fingerprints, the foregoing feature importance scores could potentially map prediction contributions onto the molecule to visualize which substructures positively or negatively impacted the prediction [25–27]. Although feature importance measures increase the explanatory power of machine learning models, caution must be taken when these scores are invoked on molecules outside the applicability domain of the model, as prediction importance does not always translate to biological relevance[28]. Locally interpretable models can be fitted to explicate predictions of black-box models. For instance, SHapley Additive exPlanations (SHAP) offers a model-agnostic method for calculating prediction-wise feature importance[21, 22]. Since this technique usually informs which features contributed to the specific test instance’s model prediction the most, it may not always lead to actionable design ideas. Thus, [24] proposed Molecular Model Agnostic Counterfactual Explanations (MMACE) to generate counterfactual explanations which would help answer the question: what changes will result in an alternate outcome, regardless of the underlying model used. These methods are based on the model’s knowledge and, therefore, may be influenced by chance correlation, rough response surfaces, and overfitted models, leading to disappointing results [29]. Recent advances in the attention mechanism of deep learners offer some explanatory power [30]. For instance, [31] uses Layer-wise Relevance Propagation to provide structural interpretation of nodes and edges (atoms and bonds), Transformer-CNN incorporates Layer-wise Relevance Propagation to calculate individual atom contributions influencing the predictions, and [32] uses salient maps to highlight the substructures closely related to the model output. These maps are analogous to the foregoing feature importance concept and have similar drawbacks in terms of deriving actionable insights for the design of new molecules.

Similarity-based methods [33] (k-nearest neighbor (KNN), kernel regression[34, 35], and pairwise kernel method[36]), provide natural intuitive interpretation at the instance-level by directly providing the training instances that influenced the model’s prediction the most. For example, read-across is a popular alternative property prediction technique that finds the most similar chemicals to the query chemical. Numerous publicly available tools use some variants of read-across techniques to aid chemists with design ideas[37]. These tools allow chemists to assess the potential of

139 the selected analogous neighbors to infer properties of the query chemical. Addition-
140 ally, similarity-based methods allow informative visualizations through network graphs
141 derived from the similarities. Network-like Similarity Graphs (NSG) [38] were devel-
142 oped to guide lead optimization in drug discovery and have often been used to display
143 the complex activity landscapes and the relationships between chemicals within a
144 target set in 2D. Expanding this to drug-target interactions, methods like Similarity
145 Ensemble Approach (SEA) [39] and Chemical Similarity Network Analysis Pulldown
146 (CSNAP) [40] enable visualization of drug-target interaction networks and the pre-
147 diction of off-target drug interactions, which have led to deeper investigations into
148 drug polypharmacology and the discovery of off-target drug interactions[41, 42]. As we
149 show later, these chemical similarity networks allow the clustering of similar molecules,
150 which enables practitioners to mine regions of desired activity for innovative design
151 ideas and potential leads. In addition to providing prediction-wise training instance
152 importance, these graph structures are directly compatible with Laplacian Scores[43]
153 for global feature importance, which have been used in QSAR modeling for feature
154 selection[44, 45]. Since SHAP and MMACE are model agnostic, they can also be paired
155 with similarity-based QSAR models to allow prediction-wise feature importance and
156 the generation of unseen counterfactuals. Thus, similarity-based methods can provide
157 multiple layers of interpretability on top of the commonly applied chemical similarity
158 interpretation and visualization methods listed above.

159 However, a problem in similarity-based QSAR is that most QSAR methods assume
160 that similar structures lead to similar activities, which is often violated in chemi-
161 cal structure modeling due to the prevalence of *activity cliffs* (ACs)[46], which are
162 pairs of compounds with similar molecular structures, but with a large difference in
163 potency against their target[47]. The existence of ACs often cause QSAR models to
164 fail, especially in the lead optimization stage [48], and limit the prediction perfor-
165 mance across the drug landscape, leading to the use of network-based methods to
166 interpret and analyze their behavior[38, 49]. One way to use similarity-based meth-
167 ods in the presence of ACs is to learn the similarity metric from the data itself,
168 instead of choosing a similarity metric *a-priori*. Large margin nearest neighbor [50] is
169 a very popular algorithm for supervised metric learning when the response variable
170 is categorical. For continuous response variables, Metric Learning Kernel Regres-
171 sion (MLKR) [51] is perhaps the most popular algorithm to estimate the similarity
172 metric. Metric learning techniques offer good explanatory power because once the
173 metric is learned, the chemical space of molecules is approximately isometric to the
174 activity space, resulting in smoother structure-activity landscapes as shown in [52].
175 Consequently, under the learned metric, high-activity molecules are clustered rela-
176 tively tightly in the chemical space and therefore, that space could be mined for
177 new molecules. Figure 1 depicts this phenomenon using various projection meth-
178 ods Generative Topographic Mapping (GTM) [53], Multidimensional Scaling (MDS),
179 t-distributed Stochastic Neighbor Embedding (t-SNE), and Uniform Manifold Approx-
180 imation and Projection (UMAP), to show the interpolated activity landscapes of the
181 protein target Coagulation factor XIII, or CHEMBL4530, in 2 dimensions and com-
182 pare them with an MLKR-based representation. Observe that, except MLKR, none
183
184

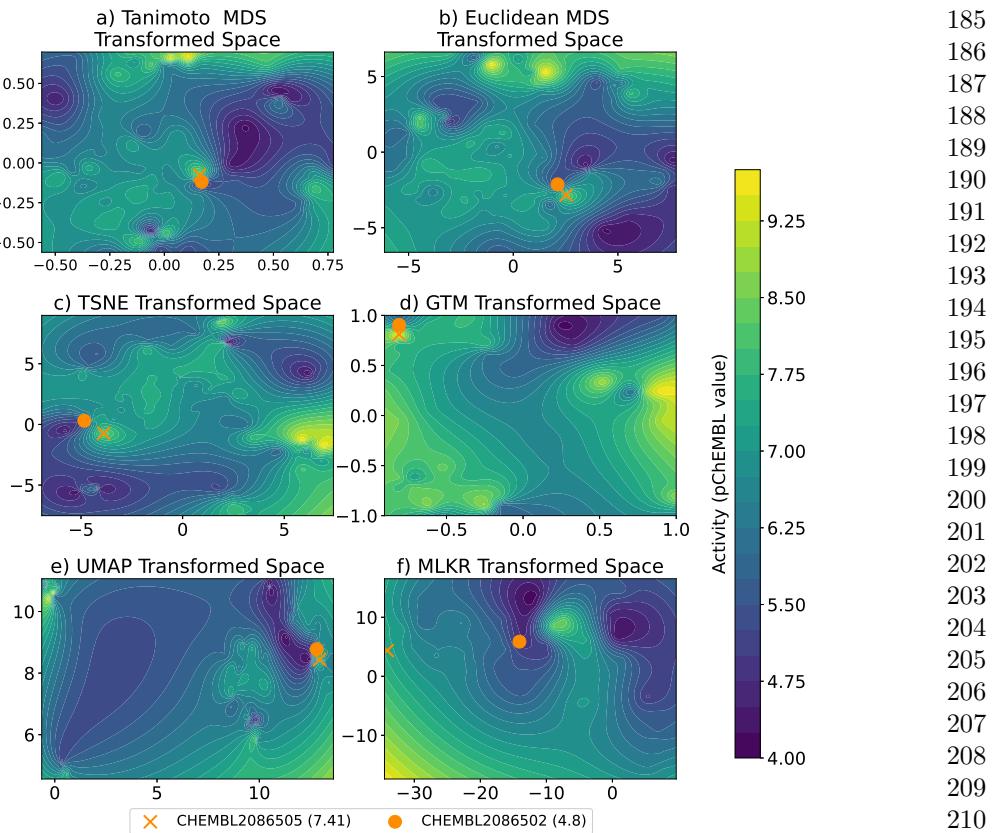


Fig. 1: The transformed chemical space and interpolated activity surface of target CHEMBL4530 using various projection methods. Transformed and interpolated chemical space using a) Multidimensional scaling (MDS) with Tanimoto distance b) MDS with Euclidian distance c) t-Distributed Stochastic Neighbor Embedding (tSNE) using Euclidian distance d) Generative Topographic Mapping (GTM) e) Uniform Manifold Approximation and Projection (UMAP) using Euclidian distance and f) Metric Learning for Kernel Regression (MLKR). Notice how activity cliffs are present regardless of the projection method. Additionally, notice that MLKR creates the smoothest activity space and best separation of the two similar (Tanimoto Similarity between ECFP4 fingerprints = 0.70) molecules, CHEMBL2086505 (pChEMBL=7.41) and CHEMBL2086502 (pChEMBL=4.8).

of the other methods were able to separate two chemically-similar-but-functionally-different molecules, CHEMBL208650 and CHEMBL2086502, which have a Tanimoto similarity between ECFP4 fingerprints of 0.70 but target difference of 2.61. This is to be expected because, in their original form, GTM, MDS, t-SNE, UMAP are all unsupervised techniques and do not incorporate the activity information in their projections. MLKR on the other hand is a supervised metric learning method, which

231 allows it to incorporate the target activity information resulting in smoother activity
232 landscapes.

233 In this paper, we develop an MLKR-inspired regression-based technique, topologi-
234 cal regression (TR), that models the distance in the response space using the distances
235 in the chemical space. TR essentially builds a parametric model to determine how
236 pairwise distances in the chemical space impact the weights of nearest neighbors in the
237 response space. Observe, unlike metric learning techniques, TR does not attempt to
238 learn a metric in the response space, nor does it attempt to provide a lower dimensional
239 projection like MDS or GTM. Rather, TR simply estimates the weights of nearest
240 neighbors. In comparison to traditional modeling methods, like RFs and SVMs, which
241 are dependent on a predefined fingerprint, TR can accommodate non-metric systems
242 and does not crucially require coordinates for each instance. As we will show in the
243 subsequent sections, TR can work on the similarities between training molecules, such
244 as those computed from molecular kernels[54, 55], thereby circumventing the problem
245 of featurization of molecules. Since, our primary use-case scenario is QSAR in the lead
246 identification/optimization process, where the contiguity of high-activity molecules
247 plays a significant role, we perform a large-scale comparison on 530 ChEMBL bio
248 targets. We use RF, ChemProp, and Transformer-CNN as baseline models and show
249 that TR matches the performance of Transformer-CNN at a significantly less com-
250 putational cost. We also observe, empirically, that TR produces numerically superior
251 predictive performance as compared to the other competing methods. Additionally,
252 both MLKR and TR produce reasonably contiguous areas of high activity, thereby
253 identifying a relatively compact high-activity chemical space.

254

255 2 Results

256

257 2.1 Model performance comparison on ChEMBL datasets

258

259 We apply our TR method with Gaussian kernel neighbor weighting on 530 ChEMBL
260 datasets under both random split and scaffold split. As explained earlier, we use the
261 ECFP4 TC distance as input to TR to predict the activity values. We use 80% of
262 all the instances in each dataset for training and the remainder for testing. For the
263 construction in section 4.2, when $I^* \cap I = \emptyset$, we use 20% of the training instances as
264 anchor points and the remaining 80% of the training set for neighborhood training.
265 We denote this method as TR* in the results. For the approach described in section
266 4.3 without disjointedness requirement, we use 50% of training instances, with a max-
267 imum of 2000 instances to improve computation time, as anchor points, and those
268 results are denoted as TR. Finally, to reduce the sensitivity of results to anchor point
269 selection, and to improve generalization error, different random sets of anchor points
270 were sampled to create an ensemble of TR models(see section 4.4). We denote this
271 method as Ensemble TR and used $t = 15$, $\mu_k = 0.6$, and $\sigma_k^2 = 0.2$ for the subsequent
272 results.

273 The average Spearman correlation and NRMSE for each method (RF, MLKR with
274 KNN, ChemProp, TCNN, TCNN with augmentation, TR*, TR, and Ensemble TR)
275 on both splitting scenarios are shown in Table 1. Figure 2 compares each method
276

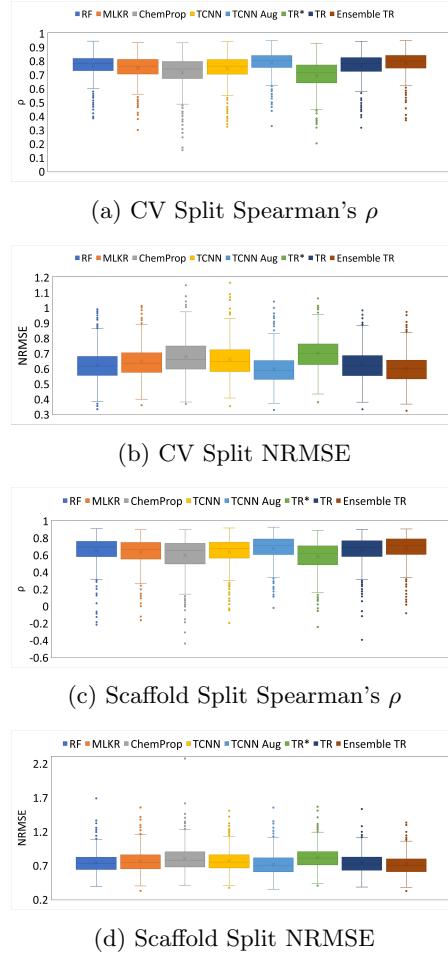


Fig. 2: Comparative analysis of model performances on the 530 ChEMBL bioactivity datasets. a) average 5-fold Cross-Validation (CV) Spearman's correlation coefficient (ρ), b) average 5-fold CV Normalized Root Mean Square Error (NRMSE) c) Scaffold split Spearman's ρ , and d) Scaffold split NRMSE. The experiment was performed on $n = 530$ ChEMBL bioactivity datasets with Random Forest (RF), Metric Learning for Kernel Regression (MLKR), ChemProp, Transformer-Convolutional Neural Network (TCNN), TCNN with augmentation (TCNN Aug), Topological Regression with disjoint anchor and training set (TR*), Topological Regression (TR), and Ensemble TR on both random cross-validation and scaffold split. The box plots show the median (central line), the interquartile range (upper and lower limits of the box), and the 5% and 95% limits (whiskers), as well as the outliers. Source data are provided in the Source Data file.

using boxplots showing the distribution of the performances for both random and scaffold splitting. As expected, TR* is unable to achieve performance comparable to the

323
 324 **Table 1: Comparative measurements obtained on**
 325 **each of the competing methods** Normalized Root
 326 Mean Square Error (NRMSE) and Spearman’s correla-
 327 tion coefficient of Random Forest (RF), Metric Learning
 328 for Kernel Regression (MLKR), ChemProp, Transformer-
 329 Convolutional Neural Network (TCNN), TCNN with
 330 augmentation (TCNN Aug), Topological Regression with
 331 disjoint anchor and training set (TR*), Topological
 332 Regression (TR), and Ensemble TR on both random
 333 cross-validation and scaffold split. Notice how both
 334 TCNN and Ensemble TR (bold) achieve numerically
 335 superior performance compared to all other competing
 336 methods.

	CV Split		Scaffold	
	Spearman	NRMSE	Spearman	NRMSE
RF	0.7629	0.6242	0.6493	0.7395
MLKR	0.7486	0.6421	0.6367	0.7593
ChemProp	0.7160	0.6776	0.5986	0.8002
TCNN	0.7437	0.6595	0.6321	0.7692
TCNN Aug	0.7858	0.5961	0.6742	0.7176
TR*	0.6935	0.7023	0.5793	0.8174
TR	0.7625	0.6255	0.6531	0.7358
Ensemble TR	0.7847	0.5989	0.6791	0.7101

346 competing methods as the model is being constrained by the disjointedness require-
 347 ment. When we relax this requirement, we observe that TR’s predictive performance
 348 improves considerably and is only numerically inferior to TCNN with augmentation.
 349 Finally, when we incorporate an ensemble of TR models, the predictive performance of
 350 Ensemble TR is essentially as good as that of TCNN with augmentation. If we invoke
 351 the law of parsimony, our conceptually straightforward, and mathematically less com-
 352 plex, topological regression approach appears to be more appealing as compared to
 353 competing deep learning techniques.

354

355 2.2 Computational comparison on ChEMBL datasets

356

357 To illustrate the computational efficiency of TR and Ensemble TR, we report each
 358 competing method’s average training time, testing time, and peak RAM consumption
 359 across all 530 datasets. These results are shown in Table 2. For fair comparison and
 360 to provide the best optimized hardware for each model, we trained the deep learning
 361 models on systems with GPUs as the training of deep learning-based models are better
 362 optimized in GPU based systems. Since the pre-trained TCNN model was released
 363 and used for fine-tuning, the reported TCNN time does not include pre-training time.
 364 From the results, we observe that TR and Ensemble TR result in the fastest training
 365 times and significantly less peak RAM consumption. For testing, TR takes more time
 366 than MLKR since RBF kernels are employed compared to MLKR which simply uses
 367 5-NN predictions after transformation, however TR still results in faster test times
 368 than TCNN. These results demonstrate the computational efficiency of TR.

Table 2: Computational complexity comparison of competing methods showing training time, testing time, and peak RAM consumption on the scaffold split. The compared methods listed on the table are Metric Learning for Kernel Regression (MLKR), Topological Regression (TR), Ensemble TR, ChemProp, Transformer-Convolutional Neural Network (TCNN), and TCNN with Augmentation (TCNN Aug).

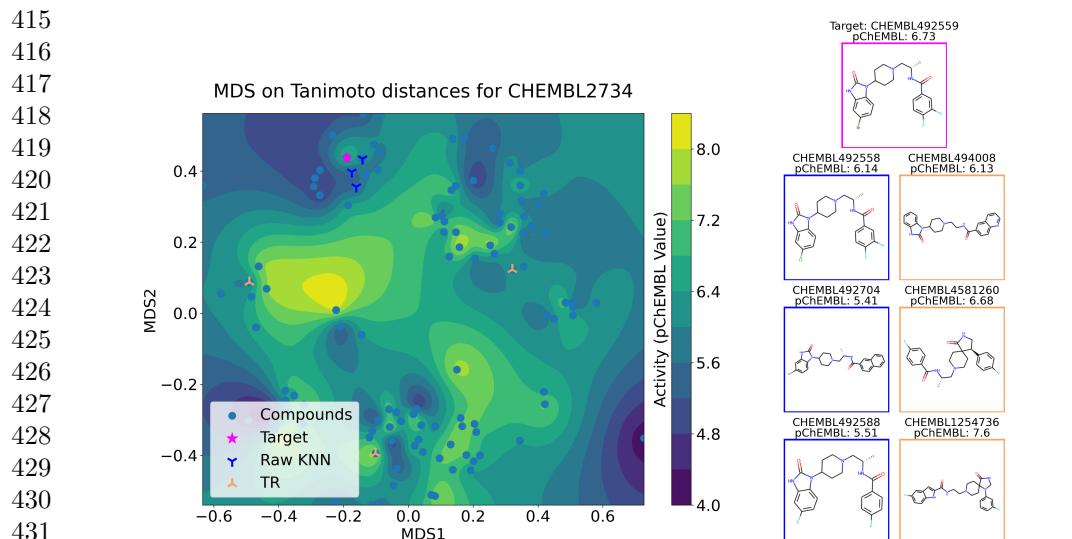
CPU System (AMD EPYC 7702, GPU System (Intel Xeon Gold 6242, 2.0 GHz, 64 Physical Cores, 128 Logical Cores, 512 GB RAM) 2.8 GHz, 16 Physical Cores, 32 Logical Cores, 384 GB RAM, NVIDIA Tesla V100)						
Method	MLKR	TR	Ensemble TR	ChemProp	TCNN	TCNN Aug
Train Time (s)	181.648	1.602	13.768	56.586	30.586	109.759
Test Time (s)	0.354	1.008	12.301	0.510	6.915	20.003
Peak RAM (GB)	1.777	0.253	0.314	2.044	2.864	3.231

2.3 Interpreting TR

Inspection of the regression coefficients in B demonstrates how TR offers more flexibility as compared to standard KNN. Recall, $W_{K,m}$, $K \in I^*, m \in I$ quantifies the impact of Y_K on Y_m . Now, in an ordinary KNN inverse distance weighting scheme, as distance between the K th instance and m th instance increases in the chemical space, $W_{K,m}$ decreases, i.e., $\frac{\delta}{\delta d_{K,m;X}^2} W_{K,m} < 0$. However, for TR $\frac{\delta}{\delta d_{K,m;X}^2} W_{K,m} = W_{K,m} b_{KK}$.

Now $W_{K,m} > 0$ by construction, therefore $\text{sign}\left(\frac{\delta}{\delta d_{K,m;X}^2} W_{K,m}\right)$ depends upon the $\text{sign}(b_{KK})$. Hence, TR can push molecules closer in chemical space far apart in the response space. What this implies is, the prediction generation process for TR can be interpreted in the same vein as that used by KNN, except, unlike KNN, TR searches for nearest set of anchor points in the response space.

We use the chemical space of the drugs targeting Phospholipase D2 (ChEMBL ID: CHEMBL2734) to demonstrate this phenomenon. In Fig. 3 we seek to predict the response corresponding to the molecule CHEMBL492559 (denoted by a red star, pChEMBL= 6.73) in the test set. Based on similarity in the chemical space, standard KNN finds three molecules, CHEMBL492558, CHEMBL492704, and CHEMBL492588, as nearest neighbors, under a 5-fold cross-validation protocol, and makes predictions based on the average of the activities of these three molecules. However, the target molecule is almost at the edge of a high-activity region. Therefore, naive KNN identifies two neighbors, CHEMBL492704 and CHEMBL492588 from the nearby low activity region (across the cliff) and only one neighbor CHEMBL492558 from the ideal high activity region. This happens because the high-activity region in the neighborhood of the target molecule is sparsely populated. In contrast, since TR directly incorporates Y in the learning, it identifies three cross-cliff molecules, CHEMBL494008, CHEMBL4581260, and CHEMBL1254736, that have greater weights in predicting the response associated with the target molecule as compared to CHEMBL492704 and CHEMBL492588. Observe that all three molecules identified by TR as nearest neighbors (CHEMBL494008, CHEMBL4581260, CHEMBL1254736) are in relatively high-activity regions. By presenting structures from diverse scaffolds that exhibit similar activities, TR not only enhances prediction



432 **Fig. 3: Comparative analysis of the neighbors found by a KNN procedure**
433 and a TR procedure. Nearest neighbors found by K-Nearest Neighbors (KNN)
434 and Topological Regression (TR) in a single fold of the 5-fold cross-validation setup
435 for the CHEMBL2734 dataset. KNN finds the nearest training samples and can lead
436 to misleading results when the target (pChEMBL=6.73) is close to an activity cliff
437 (KNN prediction=5.69), while TR attempts to find nearest neighbors in the response
438 space, leading to more informed and meaningful predictions (TR prediction=6.80).
439 The presented chemical space was done by performing a Multidimensional Scaling
440 (MDS) on the dataset. The presented molecules have a colored frame, where magenta
441 is the target presented by the same color in the MDS plot, blue are the 3 nearest
442 neighbors also presented in the MDS by the same color, and yellow is the framing the
443 molecules found by TR, presented in the MDS by the same color.

445
446 reliability but also aids in the identification of key spatial structural characteristics
447 influencing the activities. The presented structures can be further validated with
448 structural chemical methods such as structural alignment or docking simulations.

449 To further illustrate this point across the entire dataset, rather than for one
450 particular test molecule, we generated KNN-graphs depicting the predictions of the
451 various similarity-based methods with the color indicating the activity elicited by the
452 molecules. To do so, each training and test sample was represented as a node, and
453 the predicted neighbors were considered as the connecting edges. These graphs are
454 synonymous with NSGs, in fact, just like NSGs, the edges were only included if the
455 similarity was greater than a fixed cutoff TC and if the molecules were predicted as
456 one of the nearest neighbors. Therefore, the number of neighbors and the cutoff TC
457 control the connectedness of the network graphs, more connections would be estab-
458 lished with a larger number of nearest neighbors and lower cutoff similarities until the
459 graph is complete. We used 5 nearest neighbors and the mean similarity of the entire
460

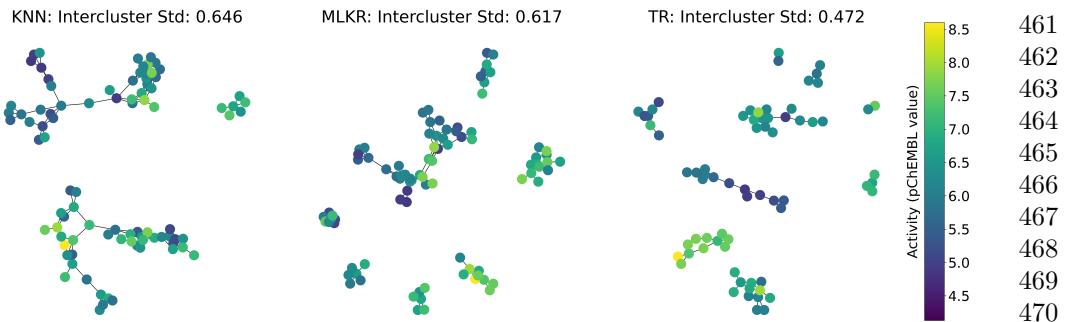


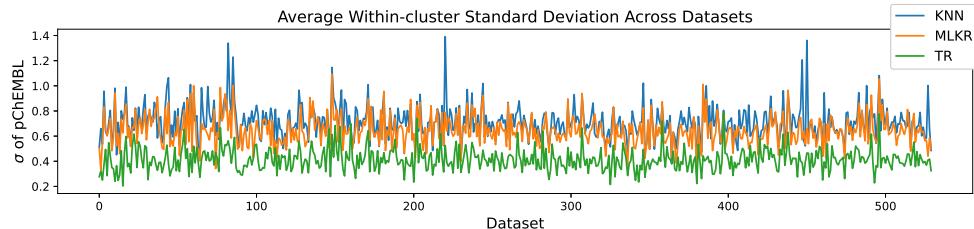
Fig. 4: k-Nearest Neighbor Graphs visualizing the 5 nearest neighbor predictions of K-Nearest Neighbors (KNN), Metric Learning for Kernel Regression (MLKR), and Topological Regression (TR). For this experiment, we used the foregoing target activity Phospholipase D2 (CHEMBL2734), and computed the 5-nearest neighbors for the three different methods, using the mean similarity of the dataset as the neighbor similarity cutoff. We can clearly see that the intercluster standard deviation is minimized by using the TR procedure.

target dataset as the cutoff TC for each competing method for all subsequent network graphs, meaning at most 5 connections would be established if their similarities were greater than the fixed cutoff TC. An example of these KNN-graphs, depicting the test nearest neighbors of a single CV fold of the dataset CHEMBL2734, is included in figure 4. Additional figures depicting the training predictions, testing predictions, and molecules within the most active cluster are included in the supplementary document. Notice that the predicted TR neighbors are similar in response value, leading to more homogeneous activity throughout the clusters, whereas KNN and MLKR both result in clusters containing diverse activity values. To quantify this variability, we included the average within-cluster standard deviations for each method in the figure where a low within-cluster standard deviation denotes a more homogeneous cluster.

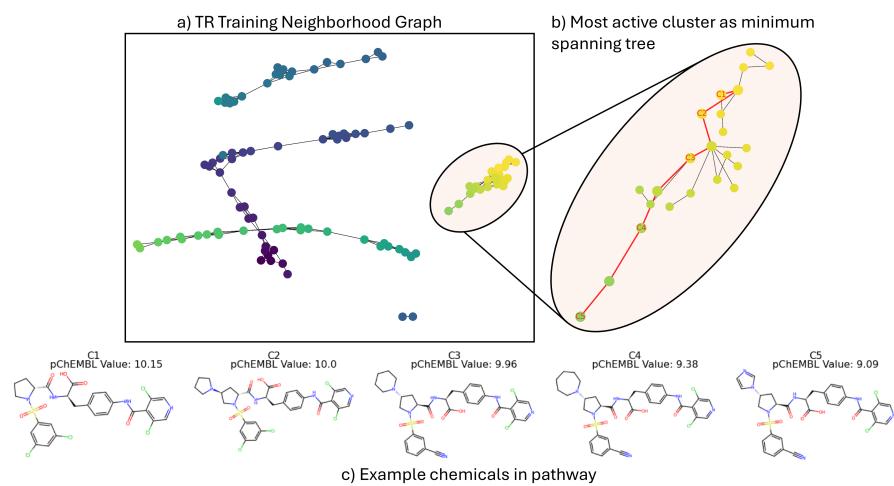
To systematically show this behavior across all 530 datasets, we calculated the average within-cluster standard deviation from the foregoing test prediction KNN-graphs for the competing methods. Figure 5 depicts these results in the form of a line graph across all 530 datasets. Clearly, TR systematically produces lower within-cluster standard deviation compared to KNN and MLKR, resulting in higher levels of homogeneous activity within the clusters. If we envision activity cliffs to be a phenomenon that induces a strong outlier within an otherwise homogeneous cluster, then it stands to reason that by measuring within-cluster homogeneity we can infer about the presence of cliffs in that cluster. Higher levels of within-cluster homogeneity essentially smooths out activity cliffs resulting in more relevant similarity-based predictions and providing practitioners with instance-wise similar molecules for lead optimization.

Since TR results in more homogeneous clusters, the clusters themselves can be more meaningfully mined by chemists for innovative design ideas, potential target leads, and lead optimization pathways. For example, clustering can be performed on the training data, and the most active cluster may contain molecules with specific

461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506



507
508
509
510
511
512
513
514 **Fig. 5: Intercluster standard deviation computed accross the 530 ChEMBL**
515 **datasets.** For this experiment, we show a quantitative comparison across all 530
516 datasets showing the average within-cluster standard deviation (σ) of pChEMBL values
517 obtained from the test prediction KNN-graphs for K-Nearest Neighbors (KNN),
518 Metric Learning for Kernel Regression (MLKR), and Topological Regression (TR). It
519 can be seen that the average standard deviation is consistently lower for TR compared
520 to the baseline methods. Source data are provided in the Source Data file.
521



522
523
524 **Fig. 6: Optimization pathway visualization in the most active training clus-**
525 **ter of target protein complex Integrin alpha-4/beta-7 (CHEMBL278).**

526 a) depicts the training neighborhood graph obtained from Topological Regression
527 (TR) predictions, b) depicts the minimum spanning tree of the most active cluster
528 with a minimum path connecting the most active and least active molecules in red,
529 and c) depicts 5 example molecules showing the lead optimization pathway.
530

531 features that practitioners can use to guide designs and future experiments. The same
532 can be done with the least active cluster to see which molecular features to avoid and
533 provide further insights. Furthermore, the most active training cluster can be mined
534 for lead molecules that have other desired characteristics, such as low toxicity or ease
535 of production. Analogous to NSGs, the training clusters can also be used to visualize
536 lead optimization pathways. Figure 6 depicts a lead optimization pathway in the most
537

active cluster of target protein complex Integrin alpha-4/beta-7 (ChEMBL278) with 553
a) the TR KNN-graph obtained from the training data of a single CV fold, b) the 554
most active cluster depicted as a minimum spanning tree with the minimum spanning 555
path between the most active and least active molecules depicted in red, and c) 556
5 example molecules from the lead optimization pathway connecting the most active 557
and least active molecules in the cluster. These pathways can be traversed by chemists 558
to envision what changes resulted in specific behaviors, allowing them to easily analyze 559
the current state of a target dataset and discover potential design ideas (additional 560
figures representing optimization pathways for various target datasets are provided in 561
the Supplementary document). If we envision an untested molecule as an additional 562
node in Figure 6(a), the TR method could directly produce the set of edges radiating 563
from that node (via the model for W) that would enable one to assess how the untested 564
molecule relates with the previously tested molecules. This could enable greater trust 565
in the predictions as the chemist could easily visualize how the new sample relates 566
to known molecules. Additionally, these graphs fit directly with Laplacian Scores for 567
feature selection, allowing global feature importance to be calculated in a routine 568
fashion. Lastly, when paired with SHAP or MMACE, which are model agnostic, TR 569
would be able to efficiently generate instance-wise feature importance and unseen 570
counterfactuals, adding additional layers to TR's interpretability. 571
572

3 Discussion

In this paper, we have developed a statistical methodology, topological regression 573
(TR), to perform similarity-based regression and demonstrated how it can be used for 574
QSAR modeling. We tested TR on regression tasks with 530 ChEMBL human targets 575
and compared it with a traditional RF, Nearest Neighbors, a metric learning algo- 576
rithm (MLKR), and two deep learning methods, ChemProp and Transformer-CNN. 577
Empirically, we observed that TR or ensemble TR compared favorably against all 578
competing methods in terms of predictive accuracy on the scaffold split and achieved 579
comparable performance with TCNN on the random splitting at a much lower com- 580
putational cost. Most importantly, TR provides explainability, visual interpretability, 581
and theoretical justifiability in the form of testable adequacy and optimal model size. 582
583

The performances of RF, TCNN, ChemProp, and MLKR are mostly interpreted in 584
a comparative sense. The usual measures employed to assess the performance of these 585
models - NRMSE, MAE - have unbounded support and hence do not offer informa- 586
tion about the goodness-of-fit. TR on the other hand completely relies on multivariate 587
general linear models - geographically weighted regression when extracting $W_{i,j}$ from 588
the drug response, and standard regression theory when modeling $W_{i,j}$. For both 589
of these techniques, rigorous tests for goodness-of-fit exist [56, 57]. Since the stan- 590
dard coefficient of determination offers an immediate goodness-of-fit statistic for linear 591
models (or transformed linear models), we compute the training R-sq values (using 592
(7)) for all 530 ChEMBL datasets considered in this paper. The average R-sq turns 593
out to be 0.8396. Evidently, our conceptually straightforward parametric linear model 594
has sufficient power to explain variation in $W_{i,j}$. Turning to predictive adequacy, we 595
compute the prediction interval for the W 's (using extracted W 's as targets) in the 596
597

599 cross-validation set. Once again, the linear model specification allows us to compute
600 the prediction interval analytically. We then compute the coverage of these prediction
601 intervals across all folds. Ideally, we would like to see the coverage of the prediction
602 interval achieve a nominal level. In all the 530 datasets across all the folds, the cover-
603 age of 95% prediction interval is 94.3%. Clearly, the model specified in (7) is adequate
604 for prediction purposes as well. These results provide empirical justification for the
605 adequacy of the TR model.

606 Given the small to moderate sample size in ChEMBL datasets, model complexity
607 has a significant impact on prediction performance. For ChemProp or TCNN-type
608 deep learners, regularization of network weights, drop-out layers, and ablations are
609 standard procedures to control model complexity. However, these measures are adhoc
610 and their theoretical properties are not well established. For standard KNN (or even
611 in MLKR), the number of neighbors determines the model size. However, we need
612 to fix the number of nearest neighbors *a-priori* and tune that quantity via cross-
613 validation. TR, on the other hand, offers a theoretically appropriate way to choose
614 neighborhood size and hence model complexity. In TR, the anchor points play the role
615 of neighbors and $|I^*|$ determines the size of the coefficient matrix B . Consequently,
616 changing $|I^*|$ yields sequences of nested models, and hence standard model selection
617 techniques, for instance, AIC or BIC, could be used to identify the appropriate size of
618 I^* without resorting to cross-validation. Since AIC/BIC automatically penalizes model
619 complexity for a given sample size, we can arrive at an optimal model complexity for
620 TR.

621 Furthermore, TR provides an intuitive explanation of its predictive mechanism
622 based on nearest neighbors in the response space as shown through KNN graphs in
623 Section 2.3. This explanation could be gleaned from MLKR as well. However, the com-
624 putational complexity associated with semi-definite programming, required in MLKR,
625 is considerable if the dimension of the input space is high. TR, on the other hand,
626 directly learns the weights associated with neighboring responses, and, by a suitable
627 transformation, estimates the parameters in an unconstrained fashion. This leads to
628 a significant reduction in computational expense as reported in Table 2.

629 Finally, the visual representation of TR’s predictive mechanism could provide
630 design ideas and allow fast knowledge-based model validation. We anticipate that our
631 framework will have practical value in drug discovery or other QSAR tasks and assist
632 in designing new molecules more effectively.

633

634 4 Methods

635

636 4.1 Data description and problem motivation

637

638 We begin with a description of the datasets that we use to illustrate the compara-
639 tive performances of the competing models. We offer a brief description of ChemProp,
640 Transformer-CNN and MLKR methods and then outline the motivation behind
641 developing the TR framework.

642 **Dataset:** Since our focus is on QSAR modeling in the lead optimization phase
643 of drug discovery, we choose to assess the performance of competing models on well-
644 curated datasets with single target bioactivity. For this purpose, we downloaded data

from the ChEMBL database[58] following the extraction protocol of [59]. This included
only selecting 'SINGLE PROTEIN' or 'PROTEIN COMPLEX' human targets with
confidence scores of 9 and 7, respectively. Additionally, only pChEMBL values, which
are comparable bioactivity measures of half-maximal response (IC50, XC50, EC50,
etc.) on a negative logarithmic scale, were selected. We refer the readers to [59] for
further data extraction details.

In the cleaning phase, we first removed the datasets that were too small to train
ChemProp and Transformer-CNN. Within each dataset, we further removed instances
with duplicated SMILES and instances with chemically invalid SMILES strings which
could not be converted to RDKit molecules. Finally, we had 530 datasets on various
human target bio-activities. Sample size ranged from 100 to 7890 with the median
sample size being 677. The various target activities, referred to as pChEMBL values,
were used as the univariate response variable.

Although several representative descriptors and fingerprints (for example: RDKit
descriptors, Mordred [6], ECFP4 [7]) are available, we mainly focus on ECFP4 rep-
resentation for similarity-based predictive models because, empirically, this represen-
tation offered the best predictive performance. We relegate the results demonstating
the superior predictive performance of the ECFP4 representation to the Supplemen-
tary Material. We calculate folded ECFP4 fingerprints using RDKit's implementation
of the Morgan algorithm with a radius of 2 atoms and bit-size of 1024. Since the out-
put of this representation system is binary, we use the Tanimoto coefficient (TC) as a
measure of similarity and $1 - TC$ as a measure of distance for TR. No standardization
steps were required as RDKit was used to extract ECFP4 fingerprints. The ECFP4
fingerprints were used to train the RF model, whereas Chemprop used the SMILES
string inputs to internally extract the graph representations and Transformer-CNN
directly used the SMILES strings.

ChemProp: We used ChemProp as a baseline model because of its demonstrated
utility in drug discovery. ChemProp is a full-fledged Graph Convolutional Neural Net-
work model that takes 2D representations of molecules as predictors. We employed
ChemProp's Bayesian hyperparameter optimization, which optimizes the hidden size,
depth, dropout, and the number of feed-forward layers, and trained the model for 100
epochs for all datasets.

Transformer-CNN: We also used Transformer-CNN (TCNN) as a baseline model
as it is self-proclaimed to be a Swiss-army knife for QSAR modeling. TCNN is a
pre-trained model on over 17 million pairs of strings for the task of SMILES canonical-
ization. The output of the transformer encoder is then used to generate model-acquired
FPs, which are used for downstream prediction through task-trained Text-CNN and
convolutional highway layers. In addition, the architecture enables data augmenta-
tion by ensembling the results from multiple non-canonical smiles for each sample.
Lastly, the architecture contains practically no hyperparameters and enables learning
rate scheduling and early stopping, limiting the need for hyperparameter optimiza-
tion. This mixture of large pre-training, sample augmentation, and string-size agnostic
architecture results in a powerful prediction model. We followed the TCNN instruc-
tions and trained the model on the SMILES strings, with and without augmentation,
for at most 35 epochs as learning rate scheduling and early stopping were employed.

691 **Metric Learning Kernel Regression:** The purpose of metric learning is to
692 find a distance metric for a specific task through supervised learning. The metric
693 found by metric learning could subsequently be used in KNN regression or kernel
694 regression for generating predictions and visualizations. For regression tasks, MLKR
695 [51] finds the Mahalanobis metric that minimizes the cumulative leave-one-out CV
696 error $\mathcal{L} = \sum_i (Y_i - \hat{Y}_i)^2$, where Y_i is the numeric response variable of the i -th training
697 sample and $\hat{Y} = \frac{\sum_{j \neq i} Y_j W_{ij}}{\sum_{j \neq i} W_{ij}}$ with $W_{..}$ being the weights associated with Gaussian
698 kernels. In particular, the transformation matrix L used to obtain the learned metric
699 can be written as a decomposition of Mahalanobis matrix $M = L^T L$. After L is
700 learned from the data, the original coordinate system of the predictor space \mathbf{X} is
701 transformed to the new coordinate system given by $L\mathbf{X}$. Thus, MLKR learns a global
702 space transformation, which can be used to calculate the distance in the response
703 space. Then KNN-regression or similarity-based kernel regression can be performed to
704 provide predictions and interpretation.

705 However, in order to compute distances, we first need to characterize the molecules
706 in a fashion such that distances can be computed. As mentioned, we focus on ECFP4
707 fingerprints, which is thus the initial coordinate system supplied to MLKR to learn the
708 transformation and produce a new coordinate system such that the predictor space is
709 approximately isometric to the response space. Fig. 7 illustrates this phenomenon. In
710 the left panel, we computed the pairwise Tanimoto distances among all the molecules
711 targeting Mitogen-Activated Protein Kinase 12 (ChEMBL ID: CHEMBL1908389)
712 using ECFP4 features and projected them in 2D MDS space. The intensity of the
713 pixels indicate the response each molecule elicited. In the right panel, we used the
714 distance metric learned from MLKR to generate the 2D coordinates. Observe how
715 the two molecules, CHEMBL3727733 and CHEMBL3729567, which appeared to be
716 neighbors in the chemical space, were pushed apart after the MLKR transformation.
717 Additionally, we observe a smoother spatial trend in the image produced after the
718 MLKR transformation which allows us to use KNN or kernel regression - with purely
719 distance-dependent kernel elements - for prediction purposes.

720 **Comparison procedure:** To compare model performances we design two types
721 of data splits: (a) random split and (b) scaffold split. Random split is done with 5-
722 fold cross-validation with 80% training and 20% testing in each fold. The scores of
723 the five folds are averaged as the final score. In drug discovery, new structures are
724 often proposed by editing on the scaffold of a known good candidate. Predictions
725 are more likely to fail across scaffolds due to greater chemical dissimilarities. Scaffold
726 split makes sure the training and test samples belong to different Murcko scaffolds
727 - mimicking scenarios when predictions for a new structure of a different scaffold is
728 sought. Since full-blown cross-validation is not feasible with scaffold splits, we use a
729 single hold-out set comprising of approximately 20% data points for each ChEMBL
730 dataset. We use Spearman ρ and Normalized Root Mean Square Error (NRMSE) to
731 compare the candidate models' capabilities to generate predictions. In Section 2, we
732 compare these two metrics obtained from ChemProp with those obtained from MLKR-
733 KNN under both splitting scenarios for all 530 ChEMBL datasets and observe that
734 MLKR-KNN offers numerically superior performance as compared to ChemProp, even
735 though MLKR is not directly a regression technique.
736

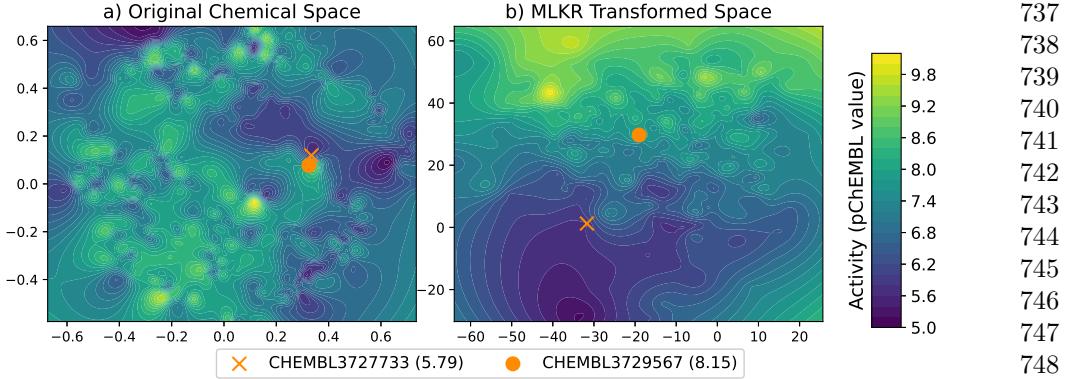


Fig. 7: The 2-D Multidimensional Scaling (MDS) of the original chemical space and Metric Learning for Kernel Regression (MLKR) transformed space of ChEMBL target Mitogen-Activated Protein Kinase 12 (CHEMBL1908389). The transformed and interpolated chemical space with a) MDS using Tanimoto distance and b) MLKR. Notice how MLKR smooths the activity space and separates the two similar molecules CHEMBL3727733 (pChEMBL=5.79) and CHEMBL3729567 (pChEMBL=8.15) compared to the original MDS transformed space.

This empirical observation motivates us to develop TR based on a distance formulation and thereby make the MLKR-type strategy amenable to statistical inference. We observe that in the MLKR procedure, a lot of effort is undertaken to ensure that the transformed space is indeed a metric space. However, for prediction, a weighted averaging of the responses from nearest neighbors is performed. Notice that symmetry and non-negativity are the only two conditions required for those weights (W_{ij}). Therefore, we contend that we can directly work with W_{ij} s instead. We then proceed to show that, under suitable distributional specification, an explicit estimator of $E(W_{ij})$ could be obtained. Since the estimand is an expectation operator, standard statistical theory (delta method, residual bootstrap) could be brought to bear to assess the statistical properties of this estimator. To the best of our knowledge, such statistical assessment of the estimates produced by vanilla MLKR is not available.

4.2 Multivariate Construction of Topological Regression

Topological regression (TR) is a similarity-based regression framework that connects the distances in the chemical space with non-negative weights appearing in nearest neighbor regression defined on the response space. The model is illustrated in Fig. 8. More specifically, we specify a multivariate regression model for the weights W_{ij} s and derive a closed-form expression for the estimator of $E(W_{ij})$ under an inverse distance-weighting scheme. Subsequently, we also offer a discussion on an approximate estimator of the foregoing quantity when the weighting is done using a Gaussian kernel.

Let \mathcal{D} represent the set of all training points. First, we partition \mathcal{D} into a set of K anchor points and $N = |\mathcal{D}| - K$ neighborhood-training points. Let $I^* = \{i_1^*, i_2^*, \dots, i_K^*\}$

783 be the set of indices associated with the anchor points and $I = \{i_1, i_2, \dots, i_N\}$ be
 784 the indices associated with the neighborhood-training points, with $I^* \cap I = \emptyset$ and
 785 $|I^*| < |I|$. Let Y_{i_j} , $i_j \in I$ be the response associated with the i_j th instance in the set
 786 I . Our goal is to express Y_{i_j} as a linear combination of responses $Y_{i_l^*}$ belonging to the
 787 set I^* , i.e.

$$788 \quad Y_{i_j} = \sum_{i_l^* \in I^*} W_{i_l^* i_j} Y_{i_l^*}, \forall i_j \in I \quad (1)$$

$$789$$

790 where $W_{i_l^* i_j}$ is a non-negative weight that determines the contribution of the
 791 response associated with the l th point in I^* towards the response associated with
 792 the j th point in I . Such non-negative weights are fairly common in distance-weighted
 793 regression, for instance, in geographically weighted spatial regression models, often
 794 the weights are specified in terms of Gaussian kernels, i.e., $W_{i_l^* i_j} = \exp(-\beta d_{i_l^* i_j}^2)$ with
 795 $d^2(\cdot)$ being a squared Euclidean distance and $\beta > 0$ controlling the smoothness of the
 796 random field.

797 **Neighborhood training model:** Customarily, the weights are expressed as a
 798 deterministic function of the distances in the predictor space. In standard KNN regression,
 799 we assume that distance in the predictor space is proportional to the distance
 800 in the response space. In metric learning, a transformation of the predictor space is
 801 learned such that there is an approximate isometry between the transformed predictor
 802 space and the response space. In TR, we instead write a formal statistical model
 803 to connect $W_{i_l^* i_j}$ with the squared Euclidean distances in the predictor space in the
 804 following fashion:

805 We define the weights

$$806 \quad W_{i_l^* i_j} = \begin{cases} 0 & \text{if } i_l^* = i_j \\ 807 & > 0 \text{ if } i_l^* \neq i_j \end{cases} \quad (2)$$

$$808$$

809 and since we have I^* and I to be disjoint and the responses could be assumed to be
 810 absolutely continuous, we can define

$$811 \quad \tilde{W}^{N \times K} = \begin{bmatrix} \log(W_{i_1^* i_1}) & \log(W_{i_2^* i_1}) & \dots & \log(W_{i_K^* i_1}) \\ 812 & \log(W_{i_1^* i_2}) & \log(W_{i_2^* i_2}) & \dots & \log(W_{i_K^* i_2}) \\ 813 & \dots & \dots & \dots & \dots \\ 814 & \log(W_{i_1^* i_N}) & \log(W_{i_2^* i_N}) & \dots & \log(W_{i_K^* i_N}) \end{bmatrix} \quad (3)$$

$$815$$

816 with the entries in \tilde{W} , i.e., $(\tilde{W})_{i_l^* i_j}$ being real quantities. Define the squared
 817 Euclidean distance matrix in the predictor space as

$$818 \quad D_X^{N \times K} = \begin{bmatrix} d_{i_1^*, i_1; X}^2 & d_{i_2^*, i_1; X}^2 & \dots & d_{i_K^*, i_1; X}^2 \\ 819 & d_{i_1^*, i_2; X}^2 & d_{i_2^*, i_2; X}^2 & \dots & d_{i_K^*, i_2; X}^2 \\ 820 & \dots & \dots & \dots & \dots \\ 821 & d_{i_1^*, i_N; X}^2 & d_{i_2^*, i_N; X}^2 & \dots & d_{i_K^*, i_N; X}^2 \end{bmatrix} \quad (4)$$

$$822$$

$$823$$

824 We define a simple multivariate linear regression model connecting \tilde{W} with D_X .
 825 Consider the m th row of \tilde{W} . Observe that, this row consists of the weights used
 826 to express the m th response in I using all the responses in I^* . We envision this
 827 row to be a set of repeated measurements taken on the m th point in I from the
 828 *vantage points* in I^* . Thus, denoting the K elements in the m th row of \tilde{W} by $\tilde{W}_{.,m} =$

$(\tilde{W}_{1,m}, \tilde{W}_{2,m}, \dots, \tilde{W}_{K,m})$, the corresponding row of predictors in D_X by $D_{.,m;X} = (d_{1,m;X}^2, d_{2,m;X}^2, \dots, d_{K,m;X}^2)$, and the matrix of regression coefficients by

$$B^{K+1 \times K+1} = \begin{bmatrix} b_{01} & b_{02} & \dots & b_{0K} \\ b_{11} & b_{12} & \dots & b_{1K} \\ \dots & \dots & \dots & \dots \\ b_{K1} & b_{K2} & \dots & b_{KK} \end{bmatrix} \quad (5)$$

we arrive at the following regression model

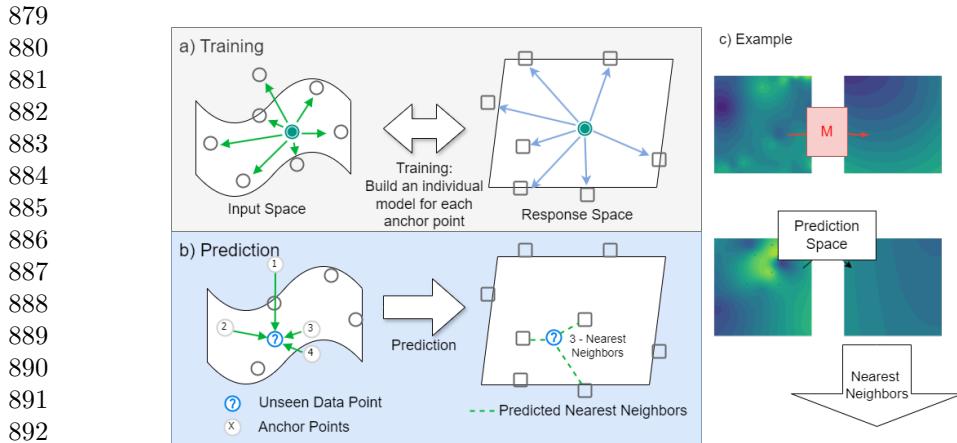
$$\begin{aligned} \tilde{W}_{1,m} &= b_{01} + b_{11}d_{1,m;X}^2 + b_{21}d_{2,m;X}^2 + \dots + b_{K1}d_{K,m;X}^2 + \epsilon_1 \\ \tilde{W}_{2,m} &= b_{02} + b_{12}d_{1,m;X}^2 + b_{22}d_{2,m;X}^2 + \dots + b_{K2}d_{K,m;X}^2 + \epsilon_2 \\ &\dots \\ \tilde{W}_{K,m} &= b_{0K} + b_{1K}d_{1,m;X}^2 + b_{2K}d_{2,m;X}^2 + \dots + b_{KK}d_{K,m;X}^2 + \epsilon_K \end{aligned} \quad (6)$$

with $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_K) \sim \mathcal{N}_K(0, \Sigma)$. Now assuming mutual independence across the N rows of \tilde{W} and since $N > K$ (by construction), we can obtain the MLEs of B and Σ . Let \hat{B} and $\hat{\Sigma}$ denote their respective estimates. Then, for a new query point, we can compute $(d_{1,query;X}^2, d_{2,query;X}^2, \dots, d_{K,query;X}^2)$ and, using \hat{B} , obtain the predictions $(\tilde{W}_{1,query}, \tilde{W}_{2,query}, \dots, \tilde{W}_{K,query})$. However, observe that (1) requires $(W_{1,query}, W_{2,query}, \dots, W_{K,query})$ to generate a prediction for the query point, and simply exponentiating the output, \hat{W} , of (6) will yield a biased estimate of W because $E(W) = E(e^{\tilde{W}}) \neq e^{E(\tilde{W})}$ due to Jensen's inequality. Therefore, we use the properties of the multivariate log-normal distribution to improve the estimate of W in the following way:

Clearly $\mathbf{W}_{.,m} = e^{\tilde{\mathbf{W}}_{.,m}}$ where the exponent is taken coordinate-wise with $\tilde{\mathbf{W}}_{.,m} \sim \mathcal{N}_K(\boldsymbol{\mu}_{.,m}, \Sigma)$ and $\mu_{j,m} = b_{0j} + b_{1j}d_{1,m}^2 + b_{2j}d_{2,m}^2 + \dots + b_{Kj}d_{K,m}^2$. Then the usual relationship between the expectation of a log-normal variate with the moment-generating function of its normal counterpart can be used to show that $E(W_{j,m}) = E(e^{\tilde{W}_{j,m}}) = \exp(\mu_{j,m} + \Sigma_{jj}/2)$. Additionally, it is fairly straightforward to show that the covariance matrix of $\mathbf{W}_{.,m}$ is given by $Var(\mathbf{W}_{.,m}) = diag(E(\mathbf{W}_{.,m}))(e^\Sigma - \mathbf{1}\mathbf{1}^T)diag(E(\mathbf{W}_{.,m}))$. Consequently, an estimator of $\mathbf{W}_{j,query}$ is given by $\hat{\mathbf{W}}_{j,m^*} = \hat{E}(\mathbf{W}_{j,query}) = \exp(\hat{\mu}_{j,query} + \hat{\Sigma}_{jj}/2)$ and the corresponding estimator of the covariance matrix is $\hat{Var}(\mathbf{W}_{.,query}) = diag(\hat{\mathbf{W}}_{.,query})(e^{\hat{\Sigma}} - \mathbf{1}\mathbf{1}^T)diag(\hat{\mathbf{W}}_{.,query})$. The estimated covariance matrix is positive definite as long as $\hat{\Sigma}$ is positive definite. Furthermore, since \hat{B} is asymptotically normally distributed, we can obtain a conservative estimate of the pointwise prediction interval of $\mathbf{W}_{.,m^*}$ using the parametric bootstrap technique outlined in [60].

Extraction of \mathbf{W} : In the above discussion, we have used $\log(\mathbf{W})$ as the target of the multivariate regression in (6). However, \mathbf{W} are not observed, but are parameters that appear in the distance-weighted regression in the response space (1). Hence, we first need to extract these weights. A naive option is to set the weights $W_{i_j^*, i_j}$ as the inverse of squared Euclidean distance in the response space between points in I and

875 I^* , i.e. $W_{i_j^*, i_j} = 1/d_{i_j^*, i_j; Y}^2$, $i_j^* \in I^*$, $i_j \in I$. In this configuration, we can simply supply
 876 $1/d_{i_j^*, i_j; Y}^2$ in the LHS of (6). We will still recover a closed form expression for $\hat{E}(W)$
 877 because the log-normal distribution is closed under an inverse transformation.
 878



893 **Fig. 8: Overview of the proposed topological regression framework.** a)
 894 Distances in the input space are used to predict distances in the response space, which
 895 can subsequently be paired with similarity-based prediction methods such as nearest
 896 neighbor or kernel regression. b) Targets of the unseen predictions will be calculated
 897 using the anchor points, and these results can be compared easily with other similarity-
 898 based methods, like the K-Nearest Neighbors (KNN). c) The whole procedure is shown
 899 in an example, where we can see the predicted space by the model (M) after training.
 900
 901
 902

903 4.3 Univariate Construction of Topological Regression

904 The requirement $I^* \cap I = \emptyset$ in the previous section induces a delicate trade-off. If
 905 we increase the number of anchor points, the neighborhood training model becomes
 906 overparametrized. If, on the other hand, we decrease the number of anchor points
 907 there may not be enough anchor points to reliably estimate the response, especially
 908 in isolated regions of high activity.
 909

910 One possible solution is to bring the distances among anchor points themselves
 911 in the neighborhood training model. But, that conflicts with the above theoretical
 912 development because each point in I^* can be observed from the remaining $K-1$ points
 913 in I^* and hence we do not have a $K \times K$ covariance matrix. Additionally, because of
 914 the symmetry constraint ($W_{i,j} = W_{j,i}$), we can only work with the triangular matrix of
 915 weights associated with points within I^* . Thus, if we forego the above multivariate log-
 916 linear regression construction (6) and view TR purely as a least-square optimization
 917 problem we can use $K(N-K)+K(K-1)/2$ equations to obtain the least-square
 918 estimates of the coefficient matrix B . In this scenario, the first $K(N-K)$ equations
 919 are obtained by varying m from 1, 2, ... N in (6). The remaining $K(K-1)/2$ equations
 920 connect the $\tilde{W}_{i_j^*, i_{j'}}$ with the instances in I^* . More specifically, dropping the subscript

i and simply denote the K elements in I^* as $\{1^*, 2^*, 3^*, \dots, K^*\}$, then we have the following system of equations:

$$\begin{aligned}
\tilde{W}_{2^*,1^*} &= b_{02} + b_{12}d_{1^*,1^*;X}^2 + b_{22}d_{2^*,1^*;X}^2 + \dots + b_{K2}d_{K^*,1^*;X}^2 + \epsilon_{2^*1^*} & 921 \\
\tilde{W}_{3^*,1^*} &= b_{03} + b_{13}d_{1^*,1^*;X}^2 + b_{23}d_{2^*,1^*;X}^2 + \dots + b_{K3}d_{K^*,1^*;X}^2 + \epsilon_{3^*1^*} & 922 \\
&\dots & 923 \\
\tilde{W}_{3^*,2^*} &= b_{03} + b_{13}d_{1^*,2^*;X}^2 + b_{23}d_{2^*,2^*;X}^2 + b_{33}d_{3^*,2^*;X}^2 + b_{K3}d_{K^*,2^*;X}^2 & 924 \\
&\dots + \epsilon_{3^*2^*} & 925 \\
\tilde{W}_{K^*,(K-1)^*} &= b_{0K} + b_{1K}d_{1^*,(K-1)^*;X}^2 + b_{2K}d_{2^*,(K-1)^*;X}^2 + b_{3K}d_{3^*,(K-1)^*;X}^2 & 926 \\
&\dots + b_{KK}d_{K^*,(K-1)^*;X}^2 + \epsilon_{K^*(K-1)^*} & 927 \\
&\dots & 928 \\
&\dots & 929 \\
&\dots & 930 \\
&\dots & 931 \\
&\dots & 932 \\
&\dots & 933 \\
&\dots & 934 \\
&\dots & 935 \\
&\dots & 936 \\
&\dots & 937 \\
&\dots & 938 \\
&\dots & 939 \\
&\dots & 940 \\
&\dots & 941 \\
&\dots & 942 \\
&\dots & 943 \\
&\dots & 944 \\
&\dots & 945 \\
&\dots & 946 \\
&\dots & 947 \\
&\dots & 948 \\
&\dots & 949 \\
&\dots & 950 \\
&\dots & 951
\end{aligned}$$

\hat{B} could be obtained by minimizing the error sum of squares. Additionally, if we assume the error terms are iid $\mathcal{N}(0, \sigma^2)$, we can easily obtain $\hat{\sigma}^2$ from the residuals. Now, when a query instance comes in with known chemical features, we can compute $\mathbf{d}_{\cdot, \text{query}}^2 = [d_{1^*, \text{query}}^2, d_{2^*, \text{query}}^2, \dots, d_{K^*, \text{query}}^2]$ in the chemical space and obtain $\hat{\mathbf{W}}_{\cdot, \text{query}} = \mathbf{d}_{\cdot, \text{query}}^2 \hat{B}$. Then an estimator of the neighborhood weights for the query point is given by $\hat{\mathbf{W}}_{\cdot, \text{query}} = \exp(\hat{\mathbf{W}}_{\cdot, \text{query}} + \hat{\sigma}^2/2)$.

Additionally, since the W 's in this case are univariate, we have the flexibility to write $W_{i_j^*, -i_j^*} = \exp(-\beta d_{i_j^*, -i_j^*; Y}^2)$ with $\beta > 0$ and replace the W 's in the LHS of (7) by $\log(d_{i_j^*, -i_j^*; Y}^2)$. Now, each d^2 has a univariate lognormal distribution. Now, to obtain an estimator of $E(W_{i_j^*, -i_j^*})$, we first observe that

$$E(W_{i_j^*, -i_j^*}) = \int_0^\infty \exp(-\beta d_{i_j^*, -i_j^*; Y}^2) f(d_{i_j^*, -i_j^*; Y}^2) dd_{i_j^*, -i_j^*; Y}^2 \quad (8)$$

is the Laplace transform of lognormal distribution. Although, there is no closed form solution of (8), but [61] derives a sharp approximator of (8) for $\beta > 0$ using *Lambert's W function*. Therefore we propose the following Monte Carlo procedure to estimate $E(W_{i_j^*, -i_j^*})$ as follows:

- a. Fit a standard geographically weighted regression with Gaussian Kernel in the response space and extract $\hat{\beta}$ [62].
- b. Fit the model (6) with $\log(d_{i_j^*, -i_j^*; Y}^2)$ in the LHS and obtain $\hat{\mu}_{i_j^*, -i_j^*}$ and $\hat{\sigma}^2$.
- c. Draw R iid replicates of $d_{i_j^*, -i_j^*; Y}^2$ from $\text{lognormal}(0, \hat{\sigma}^2)$.
- d. For each realization compute $\exp(-\hat{\beta} d_{i_j^*, -i_j^*; Y}^{2(r)} e^{\hat{\mu}_{i_j^*, -i_j^*}})$.
- e. Then the Monte Carlo estimator of the LHS of (8) is given by $\hat{E}(W_{i_j^*, -i_j^*}) = \frac{1}{R} \sum_{r=1}^R \exp(-\hat{\beta} d_{i_j^*, -i_j^*; Y}^{2(r)} e^{\hat{\mu}_{i_j^*, -i_j^*}})$

967 While this Monte Carlo approximation works well when β and σ are small, it fails to
968 explore the tail region as $\beta \rightarrow \infty$. Hence, if $\hat{\beta}$ is large, an efficient importance sampler,
969 derived in [61], should be used.

970

971 **4.4 Ensemble Topological Regression**

972 The above construction in (7) allows relaxing the disjointedness requirement $I^* \cap I = \phi$
973 to include the anchor points as neighborhood training points and allows modeling
974 the \tilde{W} 's through least squares optimization. However, by construction, $|I^*| < |I|$,
975 meaning not all training points can be included as anchor points because the least
976 squares model becomes overparameterized and overfits the training data leading to
977 poor generalization performance. Since a subset of the available training set must be
978 selected as anchor points, the results may be sensitive to the selected anchor points.
979 To average out the effect of anchor points, one can simply randomly sample multiple
980 different sets of anchor points and ensemble the results of each set. In order to achieve
981 this, we introduce Ensemble TR, which samples t sets of anchor points independently
982 and generates average predictions from the resulting t TR-models. The percentage of
983 training instances to include as anchor instances can be viewed as a hyperparameter,
984 so t percentages can be sampled from a Gaussian distribution $\mathcal{N}(\mu_k, \sigma_k^2)$, with μ_k
985 being the mean percentage of training instances to include as anchor instances and σ_k^2
986 being the requested variance of the t percentages. To verify percentage values are valid
987 and to prevent over or under-fitting, the sampled percentages are clipped between the
988 range [30%, 90%]. This leaves the user with three parameters: the number of models
989 (t), the mean percentage of training samples to include as anchor instances (μ_k),
990 and the variance of the percentages (σ_k^2). Ensemble TR maintains its computational
991 efficiency considering $D_X^{N \times N}$ can be initially calculated, and $t D_X^{N \times K}$'s can be easily
992 sampled from $D_X^{N \times N}$. This means that once distances are calculated, only t multi-task
993 linear regression models must be solved and RBF kernels applied to their outputs to
994 generate predictions, leading to fast run times.

995

996 **Data availability**

997

998 The ChEMBL datasets used in this study are available in the ChEMBL database
999 (<https://www.ebi.ac.uk/chembl/>) [58]. The code to extract the 530 ChEMBL datasets
1000 is provided in the code repository. Source data are provided with this paper.

1001

1002 **Code availability**

1003

1004 Sample data files and Python code to regenerate the TR figures and results are openly
1005 provided at https://github.com/Ribosome25/TopoReg_QSAR.

1006

1007 **References**

1008

1009 [1] Neves, B.J., Braga, R.C., Melo-Filho, C.C., Moreira-Filho, J.T., Muratov, E.N.,
1010 Andrade, C.H.: Qsar-based virtual screening: advances and applications in drug
1011 discovery. *Frontiers in pharmacology* **9**, 1275 (2018)

[2] Kwon, S., Bae, H., Jo, J., Yoon, S.: Comprehensive ensemble in qsar prediction for drug discovery. *BMC bioinformatics* **20**(1), 1–12 (2019) 1013
1014
1015

[3] Cherkasov, A., Muratov, E.N., Fourches, D., Varnek, A., Baskin, I.I., Cronin, M., Dearden, J., Gramatica, P., Martin, Y.C., Todeschini, R., *et al.*: Qsar modeling: where have you been? where are you going to? *Journal of medicinal chemistry* **57**(12), 4977–5010 (2014) 1016
1017
1018
1019
1020

[4] Grisoni, F., Ballabio, D., Todeschini, R., Consonni, V.: Molecular descriptors for structure–activity applications: a hands-on approach. *Computational Toxicology: Methods and Protocols*, 3–53 (2018) 1021
1022
1023

[5] Yap, C.W.: Padel-descriptor: An open source software to calculate molecular descriptors and fingerprints. *Journal of Computational Chemistry* **32**(7), 1466–1474 (2011) 1024
1025
1026
1027

[6] Moriwaki, H., Tian, Y.-S., Kawashita, N., Takagi, T.: Mordred: a molecular descriptor calculator. *Journal of cheminformatics* **10**(1), 1–14 (2018) 1028
1029
1030

[7] Rogers, D., Hahn, M.: Extended-connectivity fingerprints. *Journal of chemical information and modeling* **50**(5), 742–754 (2010) 1031
1032
1033

[8] Yang, K., Swanson, K., Jin, W., Coley, C., Eiden, P., Gao, H., Guzman-Perez, A., Hopper, T., Kelley, B., Mathea, M., *et al.*: Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling* **59**(8), 3370–3388 (2019) 1034
1035
1036
1037
1038

[9] Stokes, J.M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N.M., MacNair, C.R., French, S., Carfrae, L.A., Bloom-Ackermann, Z., Tran, V.M., Chiappino-Pepe, A., Badran, A.H., Andrews, I.W., Chory, E.J., Church, G.M., Brown, E.D., Jaakkola, T.S., Barzilay, R., Collins, J.J.: A deep learning approach to antibiotic discovery. *Cell* **180**(4), 688–70213 (2020) 1039
1040
1041
1042
1043

[10] Liu, G., Catacutan, D.B., Rathod, K., Swanson, K., Jin, W., Mohammed, J.C., Chiappino-Pepe, A., Syed, S.A., Fragis, M., Rachwalski, K., *et al.*: Deep learning-guided discovery of an antibiotic targeting *acinetobacter baumannii*. *Nature Chemical Biology*, 1–9 (2023) 1044
1045
1046
1047
1048

[11] Isert, C., Kromann, J.C., Stiefl, N., Schneider, G., Lewis, R.A.: Machine learning for fast, quantum mechanics-based approximation of drug lipophilicity. *ACS omega* **8**(2), 2046–2056 (2023) 1049
1050
1051
1052

[12] Wang, S., Guo, Y., Wang, Y., Sun, H., Huang, J.: Smiles-bert: large scale unsupervised pre-training for molecular property prediction. In: *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, pp. 429–436 (2019). IEEE 1053
1054
1055
1056
1057
1058

1059 [13] Karpov, P., Godin, G., Tetko, I.V.: Transformer-cnn: Swiss knife for qsar modeling
1060 and interpretation. *Journal of cheminformatics* **12**(1), 1–12 (2020)

1061

1062 [14] Doshi-Velez, F., Kim, B.: Towards a rigorous science of interpretable machine
1063 learning. *arXiv preprint arXiv:1702.08608* (2017)

1064

1065 [15] Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks.
1066 In: *International Conference on Machine Learning*, pp. 3319–3328 (2017). PMLR

1067

1068 [16] Nembrini, S., König, I.R., Wright, M.N.: The revival of the gini importance?
1069 *Bioinformatics* **34**(21), 3711–3718 (2018)

1070

1071 [17] Altmann, A., Tološi, L., Sander, O., Lengauer, T.: Permutation importance: a
1072 corrected feature importance measure. *Bioinformatics* **26**(10), 1340–1347 (2010)

1073

1074 [18] Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: Smoothgrad:
1075 removing noise by adding noise. *arXiv preprint arXiv:1706.03825* (2017)

1076

1077 [19] Koh, P.W., Liang, P.: Understanding black-box predictions via influence func-
1078 tions. In: *International Conference on Machine Learning*, pp. 1885–1894 (2017).
1079 PMLR

1080

1081 [20] Ribeiro, M.T., Singh, S., Guestrin, C.: ” why should i trust you?” explaining the
1082 predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD Interna-*
1083 *tional Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144
(2016)

1084

1085 [21] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions.
1086 *Advances in neural information processing systems* **30** (2017)

1087

1088 [22] Rodríguez-Pérez, R., Bajorath, J.: Interpretation of compound activity predic-
1089 tions from complex machine learning models using local approximations and
1090 shapley values. *Journal of Medicinal Chemistry* **63**(16), 8761–8777 (2019)

1091

1092 [23] Mothilal, R.K., Sharma, A., Tan, C.: Explaining machine learning classifiers
1093 through diverse counterfactual explanations. In: *Proceedings of the 2020 Confer-
1094 ence on Fairness, Accountability, and Transparency*, pp. 607–617 (2020)

1095

1096 [24] Wellawatte, G.P., Seshadri, A., White, A.D.: Model agnostic generation of
1097 counterfactual explanations for molecules. *Chemical science* **13**(13), 3697–3705
(2022)

1098

1099 [25] Marchese Robinson, R.L., Palczewska, A., Palczewski, J., Kidley, N.: Comparison
1100 of the predictive performance and interpretability of random forest and linear
1101 models on benchmark data sets. *Journal of chemical information and modeling*
1102 **57**(8), 1773–1792 (2017)

1103

1104

[26] Polishchuk, P.: Interpretation of quantitative structure–activity relationship models: past, present, and future. *Journal of Chemical Information and Modeling* **57**(11), 2618–2639 (2017) 1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150

[27] Balfer, J., Bajorath, J.: Visualization and interpretation of support vector machine activity predictions. *Journal of Chemical Information and Modeling* **55**(6), 1136–1147 (2015)

[28] Sheridan, R.P.: Interpretation of qsar models by coloring atoms according to changes in predicted activity: how robust is it? *Journal of chemical information and modeling* **59**(4), 1324–1337 (2019)

[29] Shoombuatong, W., Prathipati, P., Owasirikul, W., Worachartcheewan, A., Simeon, S., Anuwongcharoen, N., Wikberg, J.E., Nantasesamat, C.: Towards the revival of interpretable qsar models. *Advances in QSAR modeling: applications in pharmaceutical, chemical, food, agricultural and environmental sciences*, 3–55 (2017)

[30] Xiong, Z., Wang, D., Liu, X., Zhong, F., Wan, X., Li, X., Li, Z., Luo, X., Chen, K., Jiang, H., *et al.*: Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of medicinal chemistry* **63**(16), 8749–8760 (2019)

[31] Baldassarre, F., Azizpour, H.: Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686* (2019)

[32] Weber, J.K., Morrone, J.A., Bagchi, S., Pabon, J.D.E., Kang, S.-g., Zhang, L., Cornell, W.D.: Simplified, interpretable graph convolutional neural networks for small molecule activity prediction. *Journal of Computer-Aided Molecular Design*, 1–14 (2021)

[33] Ding, H., Takigawa, I., Mamitsuka, H., Zhu, S.: Similarity-based machine learning methods for predicting drug–target interactions: a brief review. *Briefings in bioinformatics* **15**(5), 734–747 (2014)

[34] Yamanishi, Y., Araki, M., Gutteridge, A., Honda, W., Kanehisa, M.: Prediction of drug–target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics* **24**(13), 232–240 (2008)

[35] Gajewicz-Skretna, A., Furuhamra, A., Yamamoto, H., Suzuki, N.: Generating accurate in silico predictions of acute aquatic toxicity for a range of organic chemicals: Towards similarity-based machine learning methods. *Chemosphere* **280**, 130681 (2021)

[36] Jacob, L., Vert, J.-P.: Protein-ligand interaction prediction: an improved chemogenomics approach. *Bioinformatics* **24**(19), 2149–2156 (2008)

1151 [37] Patlewicz, G., Helman, G., Pradeep, P., Shah, I.: Navigating through the mine-
1152 field of read-across tools: A review of in silico tools for grouping. Computational
1153 Toxicology **3**, 1–18 (2017)

1154

1155 [38] Wawer, M., Peltason, L., Weskamp, N., Teckentrup, A., Bajorath, J.: Structure-
1156 activity relationship anatomy by network-like similarity graphs and local
1157 structure- activity relationship indices. Journal of medicinal chemistry **51**(19),
1158 6075–6084 (2008)

1159

1160 [39] Keiser, M.J., Roth, B.L., Armbruster, B.N., Ernsberger, P., Irwin, J.J., Shoichet,
1161 B.K.: Relating protein pharmacology by ligand chemistry. Nature biotechnology
1162 **25**(2), 197–206 (2007)

1163

1164 [40] Lo, Y.-C., Senese, S., Li, C.-M., Hu, Q., Huang, Y., Damoiseaux, R., Torres, J.Z.:
1165 Large-scale chemical similarity networks for target profiling of compounds iden-
1166 tified in cell-based chemical screens. PLoS computational biology **11**(3), 1004153
1167 (2015)

1168

1169 [41] Lounkine, E., Keiser, M.J., Whitebread, S., Mikhailov, D., Hamon, J., Jenkins,
1170 J.L., Lavan, P., Weber, E., Doak, A.K., Côté, S., *et al.*: Large-scale prediction and
1171 testing of drug activity on side-effect targets. Nature **486**(7403), 361–367 (2012)

1172

1173 [42] Keiser, M.J., Setola, V., Irwin, J.J., Laggner, C., Abbas, A.I., Hufeisen, S.J.,
1174 Jensen, N.H., Kuijer, M.B., Matos, R.C., Tran, T.B., *et al.*: Predicting new
1175 molecular targets for known drugs. Nature **462**(7270), 175–181 (2009)

1176

1177 [43] He, X., Cai, D., Niyogi, P.: Laplacian score for feature selection. Advances in
1178 neural information processing systems **18** (2005)

1179

1180 [44] Sheikhpour, R., Sarram, M.A., Gharaghani, S., Chahooki, M.A.Z.: Feature selec-
1181 tion based on graph laplacian by using compounds with known and unknown
1182 activities. Journal of Chemometrics **31**(8), 2899 (2017)

1183

1184 [45] Valizade Hasanoloei, M.A., Sheikhpour, R., Sarram, M.A., Sheikhpour, E., Sharifi,
1185 H.: A combined fisher and laplacian score for feature selection in qsar based
1186 drug design using compounds with known and unknown activities. Journal of
1187 Computer-Aided Molecular Design **32**, 375–384 (2018)

1188

1189 [46] Cruz-Monteagudo, M., Medina-Franco, J.L., Pérez-Castillo, Y., Nicolotti, O.,
1190 Cordeiro, M.N.D., Borges, F.: Activity cliffs in drug discovery: Dr jekyll or mr
1191 hyde? Drug Discovery Today **19**(8), 1069–1080 (2014)

1192

1193 [47] Stumpfe, D., Hu, H., Bajorath, J.: Evolving concept of activity cliffs. ACS Omega
1194 **4**(11), 14360–14368 (2019)

1195

1196 [48] Maggiora, G.M.: On outliers and activity cliffs why qsar often disappoints. Journal
1197 of chemical information and modeling **46**(4), 1535–1535 (2006)

[49] Hu, H., Bajorath, J.: Simplified activity cliff network representations with high interpretability and immediate access to SAR information. *Journal of Computer-Aided Molecular Design* **34**(9), 943–952 (2020) 1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242

[50] Weinberger, K.Q., Blitzer, J., Saul, L.: Distance metric learning for large margin nearest neighbor classification. *Advances in Neural Information Processing Systems* **18** (2005)

[51] Weinberger, K.Q., Tesauro, G.: Metric learning for kernel regression. In: *Artificial Intelligence and Statistics*, pp. 612–619 (2007). PMLR

[52] Kireeva, N.V., Ovchinnikova, S.I., Kuznetsov, S.L., Kazennov, A.M., Tsivadze, A.Y.: Impact of distance-based metric learning on classification and visualization model performance and structure–activity landscapes. *Journal of computer-aided molecular design* **28**, 61–73 (2014)

[53] Horvath, D., Marcou, G., Varnek, A.: Generative topographic mapping approach to chemical space analysis. *Advances in QSAR Modeling: Applications in Pharmaceutical, Chemical, Food, Agricultural and Environmental Sciences*, 167–199 (2017)

[54] Fröhlich, H., Wegner, J.K., Sieker, F., Zell, A.: Kernel functions for attributed molecular graphs—a new similarity-based approach to adme prediction in classification and regression. *QSAR & Combinatorial Science* **25**(4), 317–326 (2006)

[55] Mohr, J.A., Jain, B.J., Obermayer, K.: Molecule kernels: a descriptor-and alignment-free quantitative structure–activity relationship approach. *Journal of Chemical Information and Modeling* **48**(9), 1868–1881 (2008)

[56] Charlton, M., Fotheringham, S., Brunsdon, C.: Geographically weighted regression. White paper. National Centre for Geocomputation. National University of Ireland Maynooth **2** (2009)

[57] Johnson, R.A., Wichern, D.W., et al.: Applied multivariate statistical analysis (2002)

[58] Gaulton, A., Hersey, A., Nowotka, M., Bento, A.P., Chambers, J., Mendez, D., Mutowo, P., Atkinson, F., Bellis, L.J., Cibrián-Uhalte, E., et al.: The chembl database in 2017. *Nucleic acids research* **45**(D1), 945–954 (2017)

[59] Bosc, N., Atkinson, F., Felix, E., Gaulton, A., Hersey, A., Leach, A.R.: Large scale comparison of qsar and conformal prediction methods and their applications in drug discovery. *Journal of cheminformatics* **11**(1), 1–16 (2019)

[60] Carroll, R.J., Ruppert, D.: Prediction and tolerance intervals with transformation and/or weighting. *Technometrics* **33**(2), 197–210 (1991)

1243 [61] Asmussen, S., Jensen, J.L., Rojas-Nandayapa, L.: On the laplace transform of the
1244 lognormal distribution. *Methodology and Computing in Applied Probability* **18**,
1245 441–458 (2016)

1246
1247 [62] Fotheringham, A.S., Brunsdon, C., Charlton, M.: *Geographically Weighted
1248 Regression: the Analysis of Spatially Varying Relationships*. John Wiley & Sons,
1249 (2003)

1250

1251 **Acknowledgments**

1252

1253 This work was supported in part by the National Science Foundation under
1254 Grants Nos. 2007903 (received by RP) and 2007418 (Received by S.G) and Leidos
1255 Biomed/NCI under contract 22X049 Any opinions, findings, and conclusions or recom-
1256 mendations expressed in this material are those of the authors and do not necessarily
1257 reflect the views of the National Science Foundation or Leidos Biomed/NCI. The
1258 authors acknowledge the High Performance Computing Center (HPCC) at Texas Tech
1259 University for providing computational resources that have contributed to the research
1260 results reported within this paper. URL <http://www.hpcc.ttu.edu>

1261

1262 **Author Contributions**

1263

1264 R.Z., D.N., S.G. and R.P. formulated the problem and conceived the experiments,
1265 R.Z., D.N., C.S., conducted the experiments, R.Z., D.N., C.S., S.G. and R.P. analyzed
1266 the results. All authors reviewed the manuscript. R.Z. conducted this work while he
1267 was working at Texas Tech University, however, he is currently working at Merck Inc.
1268

1269 **Competing interests**

1270

1271 The authors declare no competing interests.

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288