

Benchmarking molecular feature attribution methods with activity cliffs

José Jiménez-Luna,^{*,†,‡} Miha Skalic,[‡] and Nils Weskamp[‡]

[†]*Department of Chemistry and Applied Biosciences, RETHINK, ETH Zurich, 8093 Zurich, Switzerland*

[‡]*Department of Medicinal Chemistry, Boehringer Ingelheim Pharma GmbH & Co. KG, Birkendorfer Straße 65, 88397 Biberach an der Riss, Germany*

E-mail: joluna@ethz.ch

Abstract

Feature attribution techniques are popular choices within the explainable artificial intelligence toolbox, as they can help elucidate which parts of the provided inputs used by an underlying supervised-learning method are considered relevant for a specific prediction. In the context of molecular design, these approaches typically involve the coloring of molecular graphs, whose presentation to medicinal chemists can be useful for making a decision of which compounds to synthesize or prioritize. The consistency of the highlighted moieties alongside expert background knowledge is expected to contribute to the understanding of machine-learning models in drug design. Quantitative evaluation of such coloring approaches, however, has so far been limited to substructure identification tasks. We here present an approach that is based on maximum common substructure algorithms applied to experimentally-determined activity cliffs. Using the proposed benchmark, we found that molecule coloring approaches in conjunction with classical machine-learning models tend to outperform more modern, deep-learning-based alternatives. However, none of the tested feature attribution methods sufficiently and consistently generalized when confronted with unseen examples.

Introduction

Deep learning has quickly become a de-facto first-class citizen modeling approach in drug discovery applications, their main advantage compared to other classical machine-learning (ML) methods being their automatic feature extraction capabilities.^{1,2} Among those approaches, message-passing methods, also known as graph-neural-network models^{3,4} (GNNs) have recently become increasingly popular in chemoinformatics for relevant tasks such as molecular property prediction,⁵ de novo generative design,^{6,7} or synthesis prediction.⁸

The rise of complex deep-learning methodologies in the discussed and related fields has also been accompanied by an increasing demand of explainability, as their inner workings continue to remain elusive to interpretation among field experts.⁹ Additionally, while these models have been shown to provide impressive predictive capabilities in many use-cases, their performance, especially in fields that feature heavy experimental uncertainty, has been far from perfect, making it natural to critically assess and rationalize their predictions before decision-making. As a consequence, explainable artificial intelligence (XAI) has become a very active topic of research in theoretical ML,¹⁰ as well as within other more applied fields such as computer vision and natural language understanding.^{11,12}




Figure 1: Two example ground-truth colorings for ligand pairs present in the benchmark extracted from the 5BOT-4UM and 4F5Z-4F5 congeneric series, respectively. MCS calculations between the aforementioned pairs were carried out and non-common atoms were assigned ground-truth labels according to their activity difference sign.

In the specific context of chemoinformatics, several attempts have been made in recent years with the aim of uncovering black-box ML algorithms in property prediction tasks.^{13–15} In particular, while some studies go to great lengths to show how several modern feature attribution methods can be used to some extent to identify structural motifs,^{16,17} or property cliffs,¹⁸ it is hard to evaluate which feature attribution methods work best and under which specific conditions. Along these lines, a study by Sánchez-Lengeling *et al.*¹⁹ proposed a quantitative benchmark for several well-known feature attribution techniques in conjunction with GNNs. While it was shown that some modern feature attribution techniques can correctly highlight certain motifs, said benchmark was limited to synthetic tasks where the training procedure solely consisted in identifying whether certain molecules contained a set of predefined molecular substructures. Furthermore, non-deep-learning approaches with other classical coloring techniques were not considered in the study, while these have the advantage of working under a wider umbrella of ML models and descriptors.^{13,20–23}

In real drug discovery settings, however, one

is usually interested in the explicit prediction of pharmacologically-relevant endpoints, such as potency, or complimentary ones such as absorption, distribution, metabolism, excretion and toxicity (ADMET),²⁴ which in practice imply a certain degree of inherent experimental uncertainty.²⁵ With the goal of overcoming the limitations of previous studies, in this work we propose what we believe to be a more realistic approach to evaluate feature attribution methods for *in silico* drug discovery. We rely on maximum common substructure (MCS) algorithms to build a large data collection of biologically active pairs of closely-related compounds and use their associated activity information as a proxy for producing “ground-truth” colorings. We believe that this systematic, large-scale identification of examples yields a more relevant, comprehensive and less biased analysis than a purely qualitative validation that is based on manually-selected test cases or well-known pharmacophores extracted from the literature.¹⁸ Given their rising popularity, we evaluate several popular graph neural network architectures, as well as different associated coloring procedures, and benchmark them against other classical techniques. To our surprise, we find that a comparatively simple approach reported by Sheridan,²⁰ which uses a random forest as the underlying machine-learning model significantly outperforms all of the modern GNN-based feature attribution techniques in the proposed benchmark when the ligands present in the latter are also contained in the training sets. However, we also find out that none of the considered approaches manages to achieve satisfactory coloring performance on previously-unseen examples. Finally, we investigate the obtained results to rationalize the observed performance differences, describe possible directions for future research, and provide usage recommendations.

Materials and methods

We used two databases for different and complimentary purposes, namely the BindingDB protein-ligand validation sets^{26,27} (accessed

Jan. 2021), and the ChEMBL²⁸ database of drug-like molecules (version 27). The BindingDB protein-ligand validation sets were used as an starting point to build the proposed benchmark, which feature 1,222 molecular congeneric series of sizes ranging between 10 and 50 compounds. Having obtained an evaluation set based on closely-related compounds, additional activity data per target was necessary in order to train all underlying supervised ML models. Towards that end, we used the UniProt²⁹ identifier associated with each of the targets considered in the benchmark data and correspondingly ran a compound search in a locally-installed PostgreSQL instance of the ChEMBL database. Several selection criteria were applied: only activity information in either IC_{50} , K_d or K_i units was considered, and only training sets with at least 100 activity data points were kept. After applying these filters, 997 training sets could be successfully extracted.

Determining ground-truth colors

In order to determine ground-truth atom-level color labels for the considered benchmark sets, an implementation of the FMCS³⁰ maximum common substructure algorithm was used, as available in the rdFMCS module of the rdkit software package.³¹ MCS calculations were run for all compound pairs in each benchmark series whose activity difference exceeded 1 log units. We excluded those cases where at least one compound had a molecular weight higher than 800 Da, those pairs whose fraction of common substructure atoms was below 50%, and those whose MCS calculation time exceeded 5 minutes, for computational expense reasons.

Each identified pair represents an activity cliff and we assume that the observed potency difference has to be the result of the structural variation between the two compounds.³² Moieties present only in the more active compound are expected to receive a positive feature attribution and vice versa. Atomic labels were therefore assigned depending on the sign of the activity difference in a pair with the common substructure being considered neutral (see Fig-

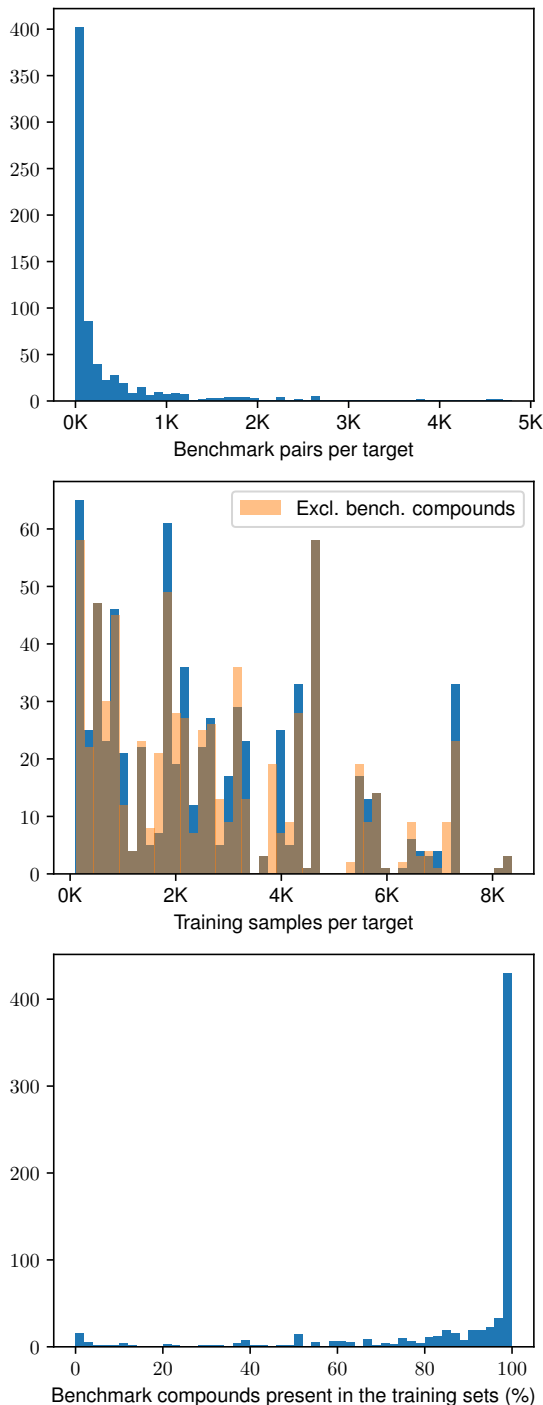


Figure 2: Histograms portraying the distribution of the number of benchmark pairs, training compounds available per target considered in the BindingDB protein-ligand validation sets (before and after removing identical ligands present in their respective benchmark sets), and percentage of benchmark ligands initially present in the training sets.

ure 1 for examples). In the benchmark analyses

described below, we assess the performance of the examined methods based on how well their coloring correlates with the ground truth.

This procedure resulted in 738 series featuring at least one colored pair of compounds, as well as available training data meeting the previously discussed criteria. Histograms describing the distribution of both the number of benchmark pairs and training compounds considered per set, as well as the percentage of benchmark compounds initially present in the training sets is provided in Figure 2.

Models

We make use of the GNN implementations provided by Sánchez-Lengeling *et al.*,¹⁹ which includes four popular variants: GraphNets,³³ Graph-Convolutional Neural Networks (GCNs),⁴ Message-Passing Neural Networks (MPNNs),³ and Graph-Attention Neural Networks (GATs).³⁴ Arguably, all of previous fall under the umbrella of message-passing algorithms, which for completeness we briefly summarize here.

Given a graph $\mathcal{G} = (V, E, u)$ with vertices $v \in V$, edges $e \in E$ and optional global graph information $u \in \mathbb{R}^k$, a graph neural network is a function f that takes a graph as an input and whose output is another graph with equal topology but with updated (*i.e.* latent) node, edge and global information. In practice, updated representations are aggregated via a read-out function into a single latent vector that can be then forward-propagated to a single scalar, so that $f : \mathcal{G} \rightarrow \mathbb{R}$. The previous architecture variants mainly differ in the choice of message function (*i.e.* the strategy which determines how node and edge updates are carried out). For all GNN block types considered, nodes and edges in the molecular graphs were featurized with the descriptors detailed in Table 1. As baselines we furthermore consider a 3-hidden-layer fully-connected neural network, and a random forest model using Extended Connectivity Fingerprints (ECFP4)³⁵ as input descriptors.

Table 1: Node and edge molecular graph features used in the training of the GNN models, as computed with the RDKit³¹ software package.

Description level	Features
Atom	atom type, chirality, valence, formal charge, hybridization, bond degree, presence in ring, aromaticity, number of hydrogens, number of radical electrons, atomic mass, van der Waals radius
Bond	bond type, bond stereo, conjugation, presence in ring

Feature attribution techniques

Several popular deep-learning feature attribution methods were used, as available in the accompanying code repository of Sánchez-Lengeling *et al.*:¹⁹ GradInput,³⁶ Class Activation Maps (CAM),³⁷ Gradient Class Activation Maps (GradCAM),³⁸ Integrated Gradients³⁹ and Attention Weights.³⁴ For completeness, a masking-like approach where the atom features in each node in the graph are sequentially zeroed-out, and its corresponding graph then forward-passed through the model, was also implemented (this method referred to as “diff” in what follows). Furthermore, for baseline purposes, the fingerprint-based masking method proposed by Sheridan²⁰ was also implemented, where the types in each atom of a molecule are sequentially changed to one that is not present in the training set, and the difference between the predictions using the unmodified and modified fingerprints is taken as a proxy for atom importance. ECFP4 fingerprints are computed for these modified molecules and then used for prediction using either a random forest, or a fully-connected neural-network model trained with the hyperparameters specified in the previous section. While there is a plethora of ML models that can also be used also in combination with other molecular fingerprinting strategies (*e.g.* Daylight fingerprints⁴⁰), we consider only the previous two combinations mainly due to their simplicity and practical popularity.




Figure 3: Predictive performance, as measured via root mean squared error (RMSE, above) and Pearson’s correlation coefficient (PCC, below) for all the 6 considered model types and 738 training sets considered in this study, using a 20% random test split in each. Dotted vertical black lines marking the median value for each model and metric combination.

Training and other details

All GNN and fully-connected-layer models were trained for a fixed number of 300 epochs, using a learning rate of 3×10^{-4} and a batch size of 32 samples. The rest of hyperparameters were set as the default ones detailed in Sánchez-Lengeling¹⁹ *et al.*. Three hidden layers with a size of 64 units were used for the node-update multilayer perceptrons in the GNN architectures, while for the fully connected models we used three-hidden layers with a size of 256 units. Random forest models were trained with 1000 base trees, and ECFP4 fingerprints with a bond radius of 2 units, computed via the rdkit software. For the Integrated Gradients feature attribution method, 500 Riemann integral approximation steps were used. Since most of the deep-learning coloring approaches produce both a score per node c_v as well as per edge $b_{u,v}$, and the proposed benchmark only considers the first, edge contributions were evenly distributed among their connecting nodes according to:

$$c'_v = c_v + \sum_{i \in \mathcal{N}(v)} \frac{b_{i,v}}{2}, \quad (1)$$

where $\mathcal{N}(v)$ is the set of neighboring vertices at one bond distance from vertex v .

Results

Model predictive performance

A summary of the predictive performance of all the considered machine-learning models (using a 20% random test-set split) and for each target considered is provided in Figure 3. Most model types show satisfactory predictive capabilities, with median Pearson correlation coefficient (PCC) values above 0.7, although the GAT model type falls slightly behind with a median PCC of 0.65. Furthermore, with a median root-mean-squared error (RMSE) and PCC values between experimental and predicted values of 0.7 and 0.82, respectively, the random forest model significantly outperforms the second best-performing alternative, namely the fully-




Figure 4: Median color accuracy and median F_1 scores (geometric average of precision and recall) at different MCS thresholds of common atom percentage for all the molecular pairs considered in the benchmark.

connected neural network that uses the same descriptor type (Wilcoxon paired signed-rank test, p -values < 0.01), and consequently the rest of the other competing GNN-based models. These results are in line with some conclusions drawn from previous related research^{41,42} where it was shown that bagging and boosting-based^{43–45} ML models performed at least on par with more modern deep-learning alternatives at a fraction of their computational cost. This is

usually the case in small-to-medium sample-size scenarios, which are arguably commonplace in hit-to-lead or lead optimization campaigns.

Molecular coloring benchmark

The main results of the proposed benchmark can be seen in Figure 4, where color accuracy and F_1 -scores are computed only taking into account the sign of the computed atomic attri-

bution, and ignoring its magnitude. To evaluate the generalization capabilities of the different XAI approaches, we report results for both the case where benchmark ligands are present during the training stage and after removing them. To investigate whether the size of substructural change between pairs has an effect on the performance of the different feature attribution methods, the proposed benchmark was further studied at different percentage thresholds of shared atoms between the considered pairs.

First looking at the case where benchmark ligands had not been removed from the training sets, it is a surprising result that the best performing molecular feature attribution approach, by a considerable margin, is the one proposed by Sheridan,²⁰ and particularly in combination with an underlying random forest model as well as ECFP4 fingerprints. Curiously, the next best-performing approach is a feed-forward neural network model using an identical featurization schema. The rest of the deep-learning-based approaches (Integrated Gradients, GradInput, CAM, GradCAM, diff) fall significantly behind these two, with performances only marginally better than random color assignment (*i.e.* 50% accuracy). With a margin of 6 absolute accuracy percent points w.r.t. the second best-performing GNN-based method (considering shared pairs with at least 50% common atoms, as per determined by the computed MCS), the exception is marked by the simple masking method (*i.e.* diff.) in combination with the GraphNet block type. Similar results can be drawn from the F_1 score plot, although in this case the latter method struggles to produce scores noticeably higher than 50% for pairs of compounds whose percentage amount of common atoms is below 70%.

Conclusions drawn from the benchmark where pairs were removed if existing in their corresponding training sets are dramatically different. None of the proposed methods, including the two best-performing ones under the previous case manage to surpass the 60% accuracy level line, and they only surpass the 50% F_1 score line only in cases of minor ligand structural differences (common atom share higher than 90%, as per determined by their MCS).

In particular, 8 of the considered combinations even failed to produce accuracies significantly higher (one-sided t -test, $\alpha = 0.01$) than random color assignment at a threshold of 50% MCS shared atoms, such as the GradInput or the Integrated Gradients approach in combination with the GCN, GraphNet and GAT block types.

This seems to either suggest that (i) many of the studied underlying ML models struggle at true mechanistic generalization (*i.e.*, the so-called Clever Hans effect⁴⁶), and can only provide meaningful explanations either if (i) the compound had been previously seen during training — a fact that is fundamentally at odds with satisfactory predictive performance as evaluated on their specific test-sets — or (ii) that current XAI techniques are unable to capture if the underlying ML models are learning activity cliffs.

Color agreement and influence of other variables

We first assessed to what degree the different molecular coloring approaches display any degree of agreement, as methods with little color correlation could be interpreted as orthogonal and could potentially provide different interpretations to specific property predictions. Towards that end, in Figure 5 we present the average Spearman’s rank correlation coefficient ρ for all coloring methods considered in the benchmark. While most approaches show a low degree of correlation, and even negative in some cases, the two fingerprint-based approaches — which coincidentally were the ones that fared better in the presented benchmark — display a moderate agreement ($\rho = 0.41$). The Class Activation Map family of methods (*i.e.* CAM and GradCAM), despite their poor performance, also show a considerable degree of agreement (about $\rho \simeq 0.5$ for some of the combinations) both within themselves and across different GNN block types (*e.g.* GCN, MPNN, GraphNet, GAT). Another interesting agreement with a similar level of correlation is the one found between the masking family of methods (*i.e.* diff.) for the GCN and MPNN block




Figure 5: Average Spearman’s ρ coefficient among the colors produced by all the feature attribution methods considered in this benchmark.

types, although these are correlated to a lesser degree ($\rho \simeq 0.3$) with the best-performing GNN block type, namely GraphNet.

We further studied whether other factors could be used to forecast the color accuracy of the different feature attribution methods. In particular, we evaluated whether the molecular similarity between the training and benchmark series, the number of training examples, and predictive performance on held-out data had an influence on attribution performance, as measured by the color accuracy reported in the previous section. In Figure 6, we present results on the influence of chemical similarity, as measured by the Tanimoto coefficient between ECFP4 fingerprints corresponding to pairs of molecules extracted from the train and bench-

mark sets, for the fingerprint-based methods and the ones used with the GCN block type. With PCC values between these two variables not noticeably higher than zero, we conclude that molecular similarity has no relevant influence on the quality of the colorings produced by any of the considered feature attribution methods. Similar conclusions can be drawn from the rest of the block types and variables considered (Figures S1-S16).

Discussion

In this work we proposed a simple benchmark based on the MCS of closely-related compounds featuring property cliffs, to evaluate the per-




Figure 6: Scatter plots portraying the average Tanimoto similarity — calculated via ECFP4 fingerprints — and color accuracy in each of the considered benchmark sets in this study. Results presented for the masking approach proposed by Sheridan²⁰ and the GNN-based feature attribution methods for the GCN block.

formance of the molecular colorings produced by feature attribution models. We believe that this benchmark represents a realistic test case that covers a wide range of compound series and is closely related to how medicinal chemists tend to think about chemical structures and how leads are typically optimized.^{47,48} In molecular machine learning, activity cliffs were in the past often falsely discarded as outliers, but are increasingly considered as particularly informative and relevant to assess the performance of a model.^{32,49}

We have furthermore shown that modern graph-based deep-learning methods, while able to correctly identify simple chemical motifs when trained on synthetically-generated datasets, struggle to correctly highlight those in real-world lead optimization datasets, even when present in the training set. This work, at some level, further highlights the importance of testing simple baselines when evaluating newly-developed approaches in molecular machine-learning research, as these were among the most performant ones. However, these conclusions need to be taken with caution, as no combination of XAI method and underlying ML model was able to successfully color previously-unseen pairs of molecules in a consistent manner (the best one only managed to do so on 56% of all tested cases).

In general, based on the results obtained

in this study, we discourage the overall use of modern feature attribution methods in prospective lead optimization applications, and particularly those that work in combination with message-passing neural networks. In general, while some methods displayed agreement with ground-truth colorings, these were only under scenarios where the colored ligands were present in the training sets. While comparatively simple ML models, such as random forests or fully-connected networks had shown the best results overall, if graph neural networks are to be used, the only combination that showed noticeably more informative results than its peers is the GraphNet block type in combination with a simple masking approach. These conclusions, however, do not imply that current feature attribution methods cannot be used for other tasks within the drug discovery pipeline. In particular, these have been proven useful in the context of structural alert identification⁵⁰ (*e.g.* identifying toxic moieties⁵¹). Alternatively, the results suggest that the machine-learning models, if providing accurate predictions, could still offer an orthogonal view to how medicinal chemists typically think of chemical structures (*e.g.* in terms of small changes to a common core)

Overall, we hope that the provided benchmark can serve as a new and more realistic starting point to evaluate explainable artificial

intelligence techniques in the context of predictive lead optimization models, and that it can motivate the development of new approaches that can successfully provide meaningful predictions in this setting.

Data and code availability

All the results presented in this study can be reproduced with the accompanying AGPLv3-licensed code repository. (https://github.com/josejimenezluna/xaibench_tf) In order to encourage the future development and testing of future molecular feature attribution methods, all pairs of compounds with their respective assigned colors are available as a compressed tarball. Instructions to download these, as well as all models and training datasets are also provided in the repository.

Acknowledgement This work was financially supported by the ETH RETHINK initiative, the Swiss National Science Foundation (grant no. 205321_182176), and Boehringer Ingelheim Pharma GmbH & Co. KG.

Conflict of interest statement. No conflict of interest declared

Supporting Information Available

- Influence of variables such as molecular similarity between training and benchmark sets, training set size, and out-of-fold performance on color agreement for all model combinations.

References

- (1) Goodfellow, I.; Bengio, Y.; Courville, A.; Bengio, Y. *Deep Learning*; MIT press Cambridge, 2016; Vol. 1.
- (2) Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Netw.* **2015**, *61*, 85–117.
- (3) Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; Dahl, G. E. Neural message passing for quantum chemistry. International Conference on Machine Learning. 2017; pp 1263–1272.
- (4) Duvenaud, D.; Maclaurin, D.; Aguilera-Iparraguirre, J.; Gómez-Bombarelli, R.; Hirzel, T.; Aspuru-Guzik, A.; Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. *arXiv preprint arXiv:1509.09292* **2015**,
- (5) Chen, H.; Engkvist, O.; Wang, Y.; Olivecrona, M.; Blaschke, T. The rise of deep learning in drug discovery. *Drug Discov. Today* **2018**, *23*, 1241–1250.
- (6) Jin, W.; Barzilay, R.; Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. International Conference on Machine Learning. 2018; pp 2323–2332.
- (7) Jin, W.; Yang, K.; Barzilay, R.; Jaakkola, T. Learning multimodal graph-to-graph translation for molecular optimization. *arXiv preprint arXiv:1812.01070* **2018**,
- (8) Coley, C. W.; Jin, W.; Rogers, L.; Jamison, T. F.; Jaakkola, T. S.; Green, W. H.; Barzilay, R.; Jensen, K. F. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chem. Sci.* **2019**, *10*, 370–377.
- (9) Jiménez-Luna, J.; Grisoni, F.; Schneider, G. Drug discovery with explainable artificial intelligence. *Nat. Mach. Intell.* **2020**, *2*, 573–584.
- (10) Adadi, A.; Berrada, M. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* **2018**, *6*, 52138–52160.
- (11) Linardatos, P.; Papastefanopoulos, V.; Kotsiantis, S. Explainable AI: A review of machine learning interpretability methods. *Entropy* **2021**, *23*, 18.

- (12) Samek, W.; Montavon, G.; Vedaldi, A.; Hansen, L. K.; Müller, K.-R. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*; Springer Nature, 2019; Vol. 11700.
- (13) Marchese Robinson, R. L.; Palczewska, A.; Palczewski, J.; Kidley, N. Comparison of the predictive performance and interpretability of random forest and linear models on benchmark data sets. *J. Chem. Inf. Model.* **2017**, *57*, 1773–1792.
- (14) Rodríguez-Pérez, R.; Bajorath, J. Interpretation of compound activity predictions from complex machine learning models using local approximations and Shapley values. *J. Med. Chem.* **2019**, *63*, 8761–8777.
- (15) Polishchuk, P. G.; Kuz'min, V. E.; Artemenko, A. G.; Muratov, E. N. Universal approach for structural interpretation of QSAR/QSPR models. *Mol. Inform.* **2013**, *32*, 843–853.
- (16) Polishchuk, P.; Tinkov, O.; Khrystova, T.; Ognichenko, L.; Kosinskaya, A.; Varnek, A.; Kuz'min, V. Structural and physico-chemical interpretation (SPCI) of QSAR models and its comparison with matched molecular pair analysis. *J. Chem. Inf. Model.* **2016**, *56*, 1455–1469.
- (17) Matveieva, M.; Polishchuk, P. Benchmarks for interpretation of QSAR models. *J. Cheminformatics* **2021**, *13*, 1–20.
- (18) Jiménez-Luna, J.; Skalic, M.; Weskamp, N.; Schneider, G. Coloring molecules with explainable artificial intelligence for preclinical relevance assessment. *J. Chem. Inf. Model.* **2021**, *61*, 1083–1094.
- (19) Sanchez-Lengeling, B.; Wei, J.; Lee, B.; Reif, E.; Wang, P.; Qian, W.; McCloskey, K.; Colwell, L.; Wiltschko, A. Evaluating attribution for graph neural networks. *Advances in Neural Information Processing Systems*. 2020; pp 5898–5910.
- (20) Sheridan, R. P. Interpretation of QSAR models by coloring atoms according to changes in predicted activity: How robust is it? *J. Chem. Inf. Model.* **2019**, *59*, 1324–1337.
- (21) Marcou, G.; Horvath, D.; Solov'Ev, V.; Arrault, A.; Vayer, P.; Varnek, A. Interpretability of SAR/QSAR models of any complexity by atomic contributions. *Mol. Inform.* **2012**, *31*, 639–642.
- (22) Riniker, S.; Landrum, G. A. Similarity maps—a visualization strategy for molecular fingerprints and machine-learning methods. *J. Cheminformatics* **2013**, *5*, 1–7.
- (23) Rosenbaum, L.; Hinselmann, G.; Jahn, A.; Zell, A. Interpreting linear support vector machine models with heat map molecule coloring. *J. Cheminformatics* **2011**, *3*, 1–12.
- (24) Van De Waterbeemd, H.; Gifford, E. ADMET in silico modelling: Towards prediction paradise? *Nat. Rev. Drug Discov.* **2003**, *2*, 192–204.
- (25) Kramer, C.; Kalliokoski, T.; Gedeck, P.; Vulpetti, A. The experimental uncertainty of heterogeneous public K_i data. *J. Med. Chem.* **2012**, *55*, 5165–5173.
- (26) 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.* **2016**, *44*, D1045–D1053.
- (27) Liu, T.; Lin, Y.; Wen, X.; Jorissen, R. N.; Gilson, M. K. BindingDB: A web-accessible database of experimentally determined protein–ligand binding affinities. *Nucleic Acids Res.* **2007**, *35*, D198–D201.
- (28) Gaulton, A.; Bellis, L. J.; Bento, A. P.; Chambers, J.; Davies, M.; Hersey, A.; Light, Y.; McGlinchey, S.; Michalovich, D.; Al-Lazikani, B., et al.

- ChEMBL: A large-scale bioactivity database for drug discovery. *Nucleic Acids Res.* **2012**, *40*, D1100–D1107.
- (29) Consortium, U. UniProt: A hub for protein information. *Nucleic Acids Res.* **2015**, *43*, D204–D212.
- (30) Dalke, A.; Hastings, J. FMCS: A novel algorithm for the multiple MCS problem. *J. Cheminformatics* **2013**, *5*, 1–1.
- (31) Landrum, G. Rdkit documentation. *Release* **2013**, *1*, 4.
- (32) Stumpfe, D.; Hu, H.; Bajorath, J. Evolving concept of activity cliffs. *ACS Omega* **2019**, *4*, 14360–14368.
- (33) Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R., et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261* **2018**,
- (34) Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903* **2017**,
- (35) Rogers, D.; Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754.
- (36) Shrikumar, A.; Greenside, P.; Shcherbina, A.; Kundaje, A. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713* **2016**,
- (37) Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; Torralba, A. Learning deep features for discriminative localization. Proceedings of the IEEE conference on computer vision and pattern recognition. 2016; pp 2921–2929.
- (38) Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. Proceedings of the IEEE international conference on computer vision. 2017; pp 618–626.
- (39) Sundararajan, M.; Taly, A.; Yan, Q. Axiomatic attribution for deep networks. International Conference on Machine Learning. 2017; pp 3319–3328.
- (40) James, C. A. Daylight theory manual. **2004**,
- (41) Sheridan, R. P.; Wang, W. M.; Liaw, A.; Ma, J.; Gifford, E. M. Extreme gradient boosting as a method for quantitative structure–activity relationships. *J. Chem. Inf. Model.* **2016**, *56*, 2353–2360.
- (42) Wu, Z.; Zhu, M.; Kang, Y.; Leung, E. L.-H.; Lei, T.; Shen, C.; Jiang, D.; Wang, Z.; Cao, D.; Hou, T. Do we need different machine learning algorithms for QSAR modeling? A comprehensive assessment of 16 machine learning algorithms on 14 QSAR data sets. *Brief. Bioinformatics* **2020**,
- (43) Friedman, J. H. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* **2001**, 1189–1232.
- (44) Chen, T.; Guestrin, C. Xgboost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. 2016; pp 785–794.
- (45) Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.-Y. LightGBM: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems* **2017**, *30*, 3146–3154.
- (46) Lapuschkin, S.; Wäldchen, S.; Binder, A.; Montavon, G.; Samek, W.; Müller, K.-R. Unmasking Clever Hans predictors and assessing what machines really learn. *Nat. Commun.* **2019**, *10*, 1–8.

- (47) Griffen, E.; Leach, A. G.; Robb, G. R.; Warner, D. J. Matched molecular pairs as a medicinal chemistry tool. *J. Med. Chem.* **2011**, *54*, 7739–7750, PMID: 21936582.
- (48) Tyrchan, C.; Evertsson, E. Matched molecular pair analysis in short: Algorithms, applications and limitations. *Comput. Struct. Biotechnol.* **2017**, *15*, 86–90.
- (49) Maggiora, G. M. On outliers and activity cliffs: Why QSAR often disappoints. *J. Chem. Inform. Model.* **2006**, *46*, 1535–1535, PMID: 16859285.
- (50) Liu, R.; Yu, X.; Wallqvist, A. Data-driven identification of structural alerts for mitigating the risk of drug-induced human liver injuries. *J. Cheminformatics* **2015**, *7*, 1–8.
- (51) Sushko, I.; Salmina, E.; Potemkin, V. A.; Poda, G.; Tetko, I. V. ToxAlerts: A web server of structural alerts for toxic chemicals and compounds with potential adverse reactions. *J. Chem. Inf. Model.* **2012**, *52*, 2310–2316, PMID: 22876798.