

Chemical Science

Accepted Manuscript

This article can be cited before page numbers have been issued, to do this please use: X. Fan, L. Guo, R. Jia, Y. Tian, Z. Yang, W. Li and B. Tian, *Chem. Sci.*, 2026, DOI: 10.1039/D6SC01813K.



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

Improving the accuracy and generalizability of molecular property regression models with a substructure-substitution-rule-informed framework

Xiaoyu Fan¹, Lin Guo¹, Ruizhen Jia¹, Yang Tian¹, Zhihao Yang¹, Weihao Li² and Boxue Tian^{1#}

¹*MOE Key Laboratory of Bioinformatics, State Key Laboratory of Molecular Oncology, Beijing Frontier Research Center for Biological Structure, School of Pharmaceutical Sciences, Tsinghua University, Beijing, 100084, China.*

²*Department of Statistics and Data Science, Tsinghua University, Beijing, 100084, China.*

To whom correspondence should be addressed:

Boxue Tian: boxuetian@mail.tsinghua.edu.cn

Abstract

Artificial Intelligence (AI)-aided drug discovery is an active research field, yet AI models often exhibit poor accuracy in regression tasks for molecular property prediction, and perform catastrophically poorly for out-of-distribution (OOD) molecules. Here, we present MolRuleLoss, a substructure-substitution-rule-informed framework that improves the accuracy and generalizability of multiple molecular property regression models (MPRMs) such as GEM and UniMol for diverse molecular property prediction tasks. MolRuleLoss incorporates partial derivative constraints for substructure substitution rules (SSRs) into an MPRM's loss function. When using GEM models for predicting lipophilicity, water solubility, and solvation-free energy (using lipophilicity, ESOL, and freeSolv datasets from MoleculeNet), the root mean squared error (RMSE) values with and without MolRuleLoss were 0.587 vs. 0.660, 0.777 vs. 0.798, and 1.252 vs. 1.877, respectively, representing 2.6–33.3% performance improvements. We show that both the number and the quality of SSRs contribute to the



magnitude of prediction accuracy gains obtained upon adding MolRuleLoss to an MPRM. MolRuleLoss yielded relative improvements on activity-cliff molecules and on a melting-point property-OOD task; however, the absolute prediction error in the OOD setting remained high, indicating that property-OOD generalization remains an open challenge. In a controlled molecular-weight extrapolation experiment, MolRuleLoss reduced the test RMSE of a GEM model from 29.507 to 0.007; because molecular weight is an exact linear function of atomic composition, this result should not be extrapolated to noisy experimental properties. We also provide a formal demonstration that the upper bound of the variation for property change of SSRs are positively correlated with an MPRM's error. Together, we show that using the MolRuleLoss framework as a bolt-on boosts the prediction accuracy and generalizability of multiple MPRMs, supporting diverse applications in areas like cheminformatics and AI-aided drug discovery.

Keywords: Substructure-substitution-rule-informed framework, Molecular property regression models, Partial derivative constraints, Activity cliff, Out-of-distribution data

Introduction

ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxicity) properties impact a compound's potential for therapeutic application.^{1,2} Compared to experimental approaches, Molecular Property Prediction Models (MPPMs) can in theory reduce costs and increase efficiency for optimizing ADMET properties.³⁻⁵ Trained humans can, based on their experience, visually evaluate the structures of compounds predicted by an MPPM, enabling selection of compounds with desired properties.⁶⁻⁸ Early MPPMs were based on (for example) empirical formulas, quantum chemistry calculations, and machine learning (ML) approaches.⁹⁻¹¹ Recent advances in deep learning (DL) have supported the development of more sophisticated MPPMs based on graph neural networks (GNNs) and Transformers.¹²⁻¹⁵ Notably, the development of these MPPMs can be fully data-driven, and these DL-based MPPMs have achieved state-of-the-art

View Article Online
DOI: 10.1039/D6SC01813K



(SOTA) performance on benchmarking datasets (*e.g.*, from MoleculeNet and MolData).

Classification models, a subtype of MPPMs, can predict a compound's status in terms of various categories (*e.g.*, whether a compound is toxic or non-toxic; whether it can pass the blood brain barrier). Molecular Property Regression Models (MPRMs), which are another subtype of MPPM that often use Simplified Molecular Input Line Entry System (SMILES) string,¹⁶ molecular graph,^{17,18} and 3D coordinates^{19,20} as input, can predict a continuous numerical value (*e.g.*, the solubility of a compound; the lipophilicity of a compound). While data scarcity is common in ADMET predictions, MPRMs have achieved high accuracy when using pre-training with large-scale datasets of unlabeled compounds followed by fine-tuning using datasets specific to desired target properties. For example, the GEM model uses both bond length and angle information for pretraining and the Uni-Mol model incorporates atomic distance information for pretraining.^{21,22}

Given that MPRMs are now integral throughout chemometrics, the ability to improve their prediction performance would be greatly welcomed. However, DL-based MPRMs suffer from overfitting on small datasets.²³ This issue manifests in two scenarios: for in-distribution (ID) data and for out-of-distribution (OOD) data.²⁴⁻²⁶ The ID overfitting scenario often relates to structurally similar molecules having large variance in properties, a situation termed "activity cliffs".^{27,28} Common solutions for addressing activity cliffs include self-supervised pretraining,^{29,30} data augmentation,^{31,32} and adding domain knowledge.^{33,34} The OOD overfitting scenario can be summarized as poor predictive accuracy of DL-based MPRMs on OOD data vs. ID data.^{35,36} This performance gap is evident across diverse DL architectures and for various molecular property prediction tasks.³⁷⁻³⁹

Despite the strong promise of MPRMs, in practice it is still common for chemists to harness their practical domain knowledge when optimizing ADMET properties in a given real-world project. A typical human domain knowledge is heuristics that predict a potential property change upon substituting one substructure with another. For example, to enhance solubility, a chemist might substitute a hydrophobic group like a phenyl with a piperazine; that is, by assuming that the $\Delta_{\text{solubility}}$ for phenyl→piperazine



will predictably increase. Analogous to these heuristics, the concept of substructure substitution rules (SSRs) in chemoinformatics is formalized through defining specific transforms (or reaction rules) that dictate how molecular substructures can be interchanged while maintaining or altering desired properties. We speculated that the prediction performance of MPRMs could potentially be improved by incorporating SSRs.

In this study, we introduce MolRuleLoss, a substructure-substitution-rule-informed framework designed to enhance the accuracy and generalizability of MPRMs by incorporating constraints for SSRs into an MPRM's loss function. Recent molecular representation learning methods also incorporate inter-molecular substructure variation, such as MolCLR⁴⁰ and MolMCL⁴¹. In contrast to these representation-level contrastive frameworks, MolRuleLoss encodes MMPA-derived quantitative rules as derivative constraints directly within the supervised objective, offering a complementary formulation. MolRuleLoss boosts the accuracy of state-of-the-art (SOTA) MPRMs (GEM and UniMol) for diverse molecular property prediction tasks, including predicting lipophilicity, water solubility and solvation-free energy. We show that both the number and quality of rules contribute to the magnitude of prediction accuracy gains obtained upon adding MolRuleLoss to an MPRM. Moreover, MolRuleLoss improves the generalizability of MPRMs for "activity cliff" and OOD molecules. In an extreme scenario for molecular weight prediction addressing the OOD problem, MolRuleLoss reduces the root mean squared error (RMSE) of a GEM model from 29.507 to 0.007. Finally, upon considering the prediction performance gain data across the multiple tasks we examined in this study, we discovered an informative and practically useful trend—that model error increases along with increases in the STD_{\max} value (*i.e.*, property change variation)—which we formalized as our rule-variation error bound observation with a mathematical proof. Our study shows that adding the MolRuleLoss framework enables more accurate MPRM predictions without changing the original MPRM's architecture, like a bolt-on to support diverse MPRMs for a wide range of molecular property prediction applications in cheminformatics.



RESULTS

MolRuleLoss adds substructure substitution rules from the MMPA method to the loss function of MPRMs

Previous studies of molecular property prediction have shown that incorporating knowledge can improve predictive accuracy.^{8,42} In many MPRMs, a standard mean squared error (MSE) term that captures the difference between predicted and true labels is often used as a loss function, providing a signal for optimizing model parameters.⁴³⁻⁴⁵ Inspired by Physics-Informed Neural Networks (PINNs),^{46,47} which embed physical formulas into DL models, we speculated that using partial derivatives to impose constraints based on a set of substructure substitution rules for a specific task (*e.g.*, predicting lipophilicity and solubility) would potentially improve MPRM accuracy. Pursuing this, we designed the MolRuleLoss framework, leveraging a set of SSRs derived from the matched molecular pair analysis (MMPA) method⁴⁸ seeking to enhance the accuracy of various MPRMs, ideally for multiple prediction tasks. Our MolRuleLoss framework allows for the automated extraction of SSRs directly from chemical datasets, which can subsequently be applied as constraints to an MPRM.

MolRuleLoss comprises an MSE term and an SSR term for quantifying property changes for the replacement of one molecular substructure with another (Fig. 1a). Lambda (λ) weighs the SSR term relative to the MSE term in the loss function (Eqn. 1). Notably, our experience ultimately showed that λ should be adjusted according to the specific desired prediction task. The default value of λ in MolRuleLoss is 0.3. When implementing an SSR in the loss function, we relate the output value of a DL model, y , to the change in the number of x_i and x_j substructures (see Eqns. 3-5 in Methods). The n-D vector saves the number of substructures for each molecule, which is converted into SSRs automatically within the MolRuleLoss framework (Fig. 1a).

$$Loss = Loss_{MSE} + \lambda \times Loss_{SSR} \quad (Eqn. 1)$$

The specific SSRs we implemented in MolRuleLoss are based on pairwise



comparison of two molecules in the training dataset (that is, excluded from the test set, seeking to prevent data leaking); these two molecules are identical except for a single substructure substitution. For instance, with a common molecular scaffold “R”, one molecule might have “A” attached (R-A) and another “B” (R-B) (Fig. 1b). The structural change from A to B induces a property change ($\Delta P = P_{[R-A]} - P_{[R-B]}$), like solubility or lipophilicity (logP), which yields a quantifiable rule for the substitution of substructure A with substructure B that corresponds to a property change (ΔP).

Considering that the property changes caused by substructure transformations vary across different molecules, we introduce an adaptive unit. For each rule, a learnable parameter vector with the same dimensionality as the molecular features is added. The dot product between the molecular feature vector and this learnable parameter vector serves as the adaptive unit, which is incorporated into the SSR term in the loss function. Mathematically, such SSRs can be modeled using partial derivatives, where the derivative of the molecular property with respect to the count of substructure A, minus the derivative with respect to substructure B, equals ΔP . Statistical insights, such as the average property change and its variability (defined by the STD) can be calculated (Fig. 1a). MolRuleLoss is a framework that can be applied as a bolt-on to diverse MPRMs.

View Article Online
DOI: 10.1039/D6SC01813K



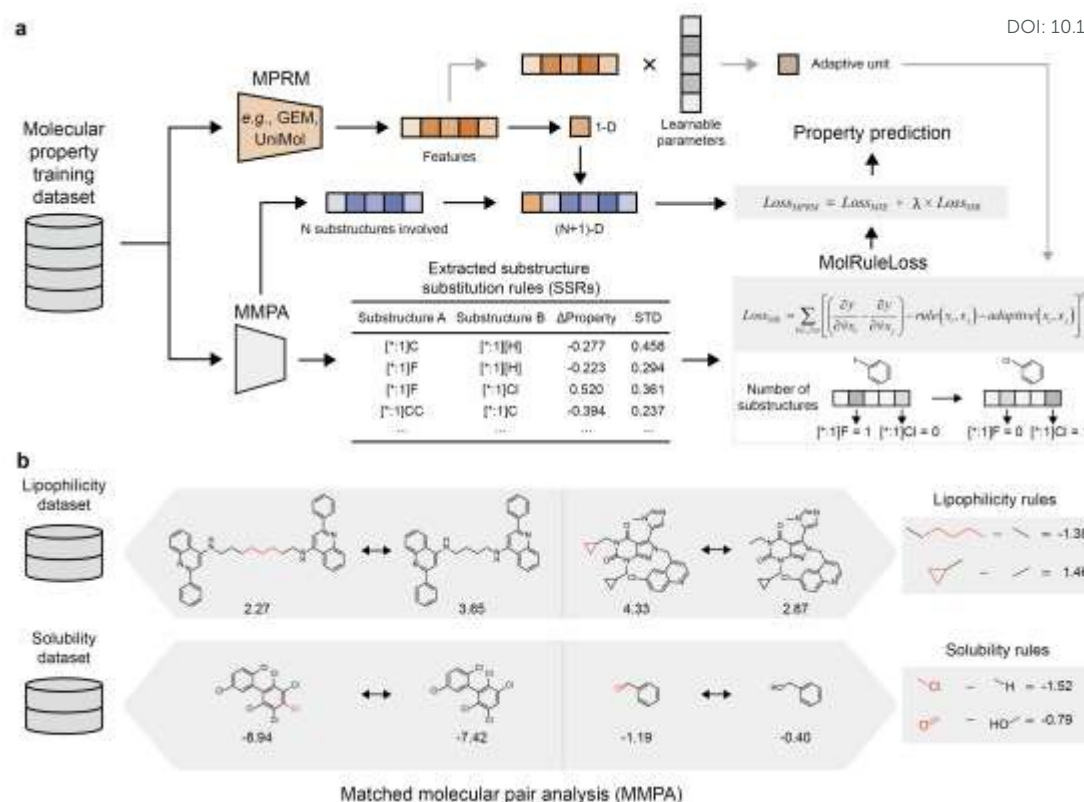


Figure 1 | Overview of the MolRuleLoss Framework. Our MolRuleLoss framework leverages a set of substructure substitution rules (SSRs) derived from the matched molecular pair analysis (MMPA) method, seeking to enhance the accuracy of MPRMs for multiple prediction tasks. **a**, MolRuleLoss comprises an MSE term and an SSR term for quantifying the property changes of SSRs. λ controls the relative importance of SSRs for MolRuleLoss's MSE term (Eqn. 1). When implementing an SSR in the loss function of an MPRM, we relate the output value of a DL model, y , to the change in the number of x_i and x_j substructures (see Eqns. 3-5 in Methods). Extracted SSRs will be converted to rules automatically within the MolRuleLoss framework. **b**, Examples of SSRs. The specific SSRs from the MMPA method are based on pairwise comparison of two molecules in the training dataset; these two molecules are identical except for a single substructure substitution. The structural change from A to B induces a property change ($\Delta P = P_{[R-A]} - P_{[R-B]}$), like solubility or lipophilicity (logP).

Improving the accuracy of multiple MPRMs for diverse tasks using MolRuleLoss

We initially applied MolRuleLoss to MPRMs for the lipophilicity prediction task, using the MoleculeNet lipophilicity dataset (Lipo), which contains 4,200 entries: each entry



is a molecule structure in SMILES format and its corresponding logP value. We divided this dataset into training, validation, and testing sets at an 8:1:1 ratio using the same scaffold split method employed for the development of previously reported MPRMs.^{49,50} We used GEM (Geometry-Enhanced Molecular Representation Learning) as our base model; GEM uses a purpose-built geometry-based graph neural network architecture and a set of dedicated geometry-level self-supervised learning strategies to capture molecular geometry knowledge.¹⁷ Note that the results we report are the average over five random seeds. For the test set, the RMSE of the GEM model with MolRuleLoss was 0.587, compared to 0.660 without it, indicating that MolRuleLoss improves the predictive accuracy of the GEM model (Fig. 2a).

We next trained UniMol and GEM models for predicting lipophilicity, water solubility and solvation-free energy (using the lipophilicity, ESOL, and freeSolv datasets from MoleculeNet), all with and without MolRuleLoss (Fig. 2b-c). For the GEM model on the three tasks, the RMSE values with and without MolRuleLoss were 0.587 ± 0.014 vs. 0.660 ± 0.010 , 0.777 ± 0.018 vs. 0.798 ± 0.020 , and 1.252 ± 0.130 vs. 1.877 ± 0.150 , representing 2.6% to 33.3% performance improvements ($P < 0.001$ for the lipophilicity and freeSolv tasks; $P < 0.05$ for the ESOL task, Fig. 2b). Note that Uni-Mol is a pretrained model with the SE(3) Transformer architecture that achieved SOTA performance in multiple MPPM contests.¹⁹ For the three tasks with the UniMol model, the RMSE values with and without MolRuleLoss were 0.555 vs. 0.603, 0.773 vs. 0.788, and 1.293 vs. 1.480 ($P < 0.001$ for all three tasks, 1.9% to 12.6% performance improvements; Fig. 2c). These results show that MolRuleLoss boosts the accuracy of multiple MPRMs for diverse molecular property prediction tasks. Importantly, using a randomly initialized GIN backbone without pretraining, MolRuleLoss reduced the lipophilicity test RMSE from 0.721 ± 0.016 to 0.686 ± 0.005 (paired t-test $P = 0.015$), indicating that the observed improvement is not contingent on GEM pretraining (Table S1). It bears emphasis that, given that UniMol and GEM are highly optimized MPPMs with recognized SOTA performance, the increases gained by adding MolRuleLoss would be expected to manifest as substantive

View Article Online
DOI: 10.1039/D6SC01813K



improvements for the set of downstream applications (for example in ADMET optimizations) that use the property prediction values as their foundation.

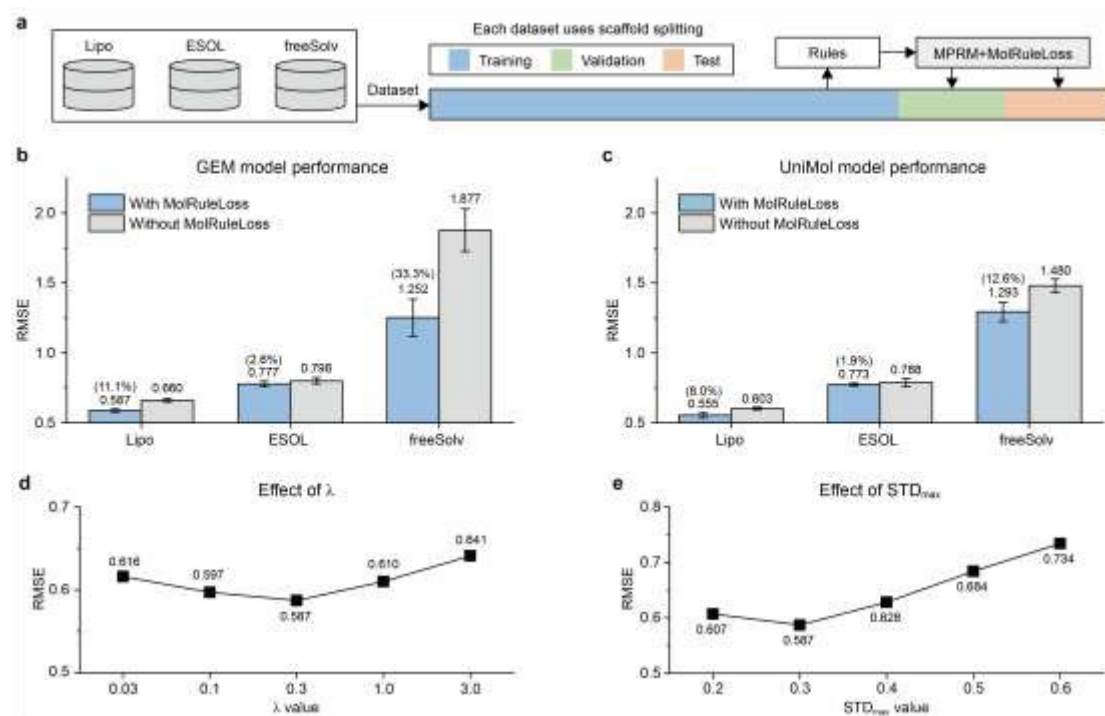


Figure 2 | Comparison of MPRMs with and without MolRuleLoss. **a**, The data splitting scheme and SSR extraction process. **b**, RMSE comparison for GEM models with or without MolRuleLoss, using the lipophilicity, ESOL, and freeSolv datasets from MoleculeNet. **c**, UniMol models were trained for predicting the same tasks, with and without MolRuleLoss. GEM and UniMol are highly optimized MPPMs with state-of-the-art (SOTA) performance on MoleculeNet; our results show that MolRuleLoss boost the accuracy of multiple MPRMs for diverse property prediction tasks. **d-e**, Assessing the influence of MolRuleLoss hyperparameters on lipophilicity prediction performance. **d**, The lipophilicity prediction accuracy of GEM using the indicated distinct λ values; λ weighs the SSR term relative to the MSE term in the loss function. **e**, The lipophilicity prediction accuracy of the GEM models using the indicated STD_{max} values.

Assessing the influence of MolRuleLoss hyperparameters on prediction accuracy

The accuracy of MPRMs using our MolRuleLoss framework are in theory affected by two hyperparameters: the Lambda value (λ), which weighs the SSR term relative to the MSE term in the loss function (Eqn. 1); and the STD threshold used to decide which



SSRs to include (STD_{max}), indicating the property variation for the SSR when it alters properties across various molecular scaffolds. We used GEM as the base model for evaluation of hyperparameter influence of MolRuleLoss on the lipophilicity prediction task (using the Lipo dataset). Our analysis started with $\lambda=0.3$ and $STD_{max}=0.3$, and we only retained SSRs with an occurrence higher than 10 (Methods). We found that $\lambda=0.3$ performed best for the lipophilicity prediction task, with an RMSE value of 0.587 on the test set; the RMSEs for $\lambda=3.0$ and $\lambda=0.03$ were 0.641 and 0.616 (Fig. 2d). These results establish that the accuracy of an MPRM using MolRuleLoss is sensitive to λ .

Given that not all molecules in the testing set possess substructure transformations defined by our SSR set, we expect that the number and quality of rules (defined by STD_{max} , with a low value indicating high quality) will affect the accuracy of MPRMs with MolRuleLoss. To this end, we evaluated the influence of STD_{max} (at $\lambda=0.3$) with GEM as the base model and using the Lipo dataset. The RMSE values were 0.607 and 0.734 for $STD_{max}=0.2$ and $STD_{max}=0.6$ (Fig. 2e). When STD_{max} is small, fewer rules are obtained; when STD_{max} is large, low quality rules are obtained. These results show that both the number and quality of rules affect the accuracy of MPRMs. The adaptive unit in MolRuleLoss introduces a single shared learnable matrix $\theta \in \mathbb{R}^{\{|V| \times d\}}$ ($\approx 30 \times 32 \approx 10^3$ parameters, $\sim 1\%$ of the GEM backbone); a parameter-matched PM-GEM control yields nearly identical RMSE to GEM (0.658 ± 0.012 vs. 0.660 ± 0.010), indicating that the observed improvement arises from the rule-based constraint rather than additional model capacity (Table S2).

The influence of MolRuleLoss training data dimensionality and rule number on prediction accuracy


Previous studies have established that MPRM accuracy can be improved by increasing dataset volume.³⁸ In theory, an MPRM can automatically learn “rules” of a task in a “big data” regime (Fig. 3a). To investigate the influence training data dimensionality on the prediction accuracy upon adding MolRuleLoss to an MPRM, we applied MolRuleLoss to a GEM model for predicting the melting point (MP) of compounds.



We used an MP dataset comprising 275K entries (each entry comprising a SMILES structure and a temperature value) from OChem.⁵¹ Based on this MP275K dataset, we generated a test set of 175K entries using scaffold splitting, and then extracted SSRs from the training dataset to train MPRMs for accuracy evaluation on the same MP175K test set (Fig. 3b).

We first trained GEM models without MolRuleLoss at different data volumes, including 0.5K, 5K, 50K, and 100K (Fig. 3c). The RMSE values for GEM without MolRuleLoss on the 175K test set were, respectively, 51.372, 46.700, 39.705, and 38.291. These results show that expanding training data dimensionality boosts the accuracy. We then trained GEM models with MolRuleLoss on the 175K test set, and the RMSE values for GEM with MolRuleLoss were 49.082, 43.290, 38.748, and 37.422 ($\lambda=0.3$, $\text{STD}_{\text{max}}=0.27$) (Fig. 3c). These results show that MolRuleLoss consistently boosts the accuracy of GEM, even upon massively expanding the dimensionality of training data.

Considering that the SSR set derived from a small training dataset is relatively sparse (the number of SSRs for 0.5K is 44, while the number of SSRs for 100K is 175), we speculated that adding more SSRs from another source, *e.g.*, from another dataset for the same property, might improve the prediction accuracy of MPRMs. This is conceptually analogous to “transfer learning”.⁵² To test this, we transferred 175 SSRs obtained with MolRuleLoss upon training with the MP100K dataset to an MP0.5K dataset. For the MP0.5K dataset, the GEM model without MolRuleLoss yielded an RMSE of 51.372. When 44 SSRs extracted from MP0.5K were applied through MolRuleLoss, the RMSE improved by 4.4% to 49.082. Using 175 MP100K SSRs as constraints led to an even greater improvement, reducing the RMSE to 46.948, an 8.6% improvement over the GEM without MolRuleLoss (Fig. 3e). These results indicate that expanding the number of rules can improve the accuracy of MPRMs with MolRuleLoss. Notably, while data in drug discovery is generally expensive and time-consuming to obtain, our results show that using rules can apparently help address the data scarcity issue in ADMET optimization efforts.



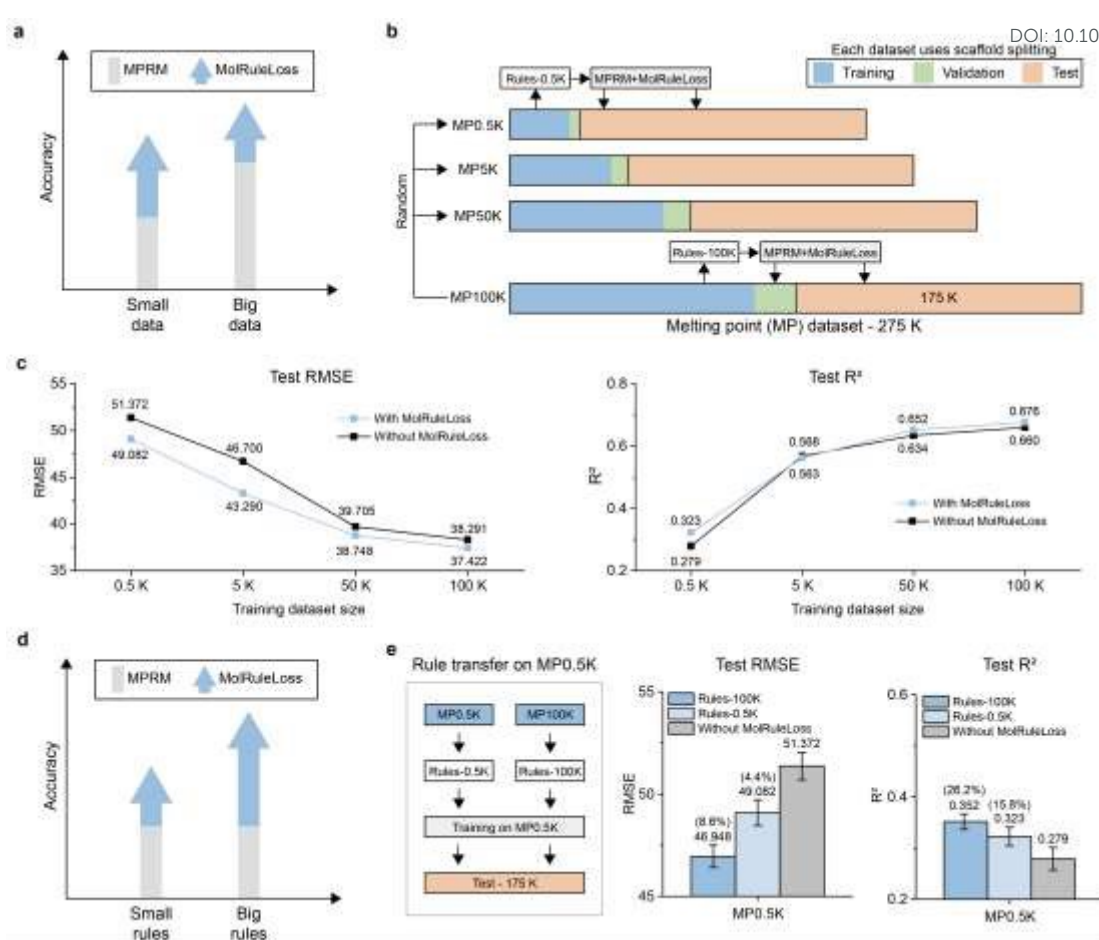



Figure 3 | Impact of training data and rule number on MPRM performance when using MolRuleLoss. **a**, We expect that MolRuleLoss would confer relatively small gains in molecular property prediction accuracy upon expanding the dimensionality of the training data. **b**, The data splitting and SSR extraction process used for studying whether MolRuleLoss improves MPRM accuracy at large training dataset size using the melting point (MP) dataset. The MP dataset contains 275K entries (with each entry comprising a SMILES structure and a temperature value). We trained GEM models with and without MolRuleLoss under different data volumes (0.5K, 5K, 50K, and 100K). **c**, Accuracy comparisons based on RMSE and R^2 . **d**, Transferring rules from a large dataset with MolRuleLoss in theory improves accuracy of MPRM, which is analogous to the “transfer learning” concept. **e**, RMSE comparison for the “no rules”, MP0.5K-rules, and MP100K-rules test iterations of the GEM model using the MP dataset.

Including low standard deviation rules with MolRuleLoss boosts the accuracy and generalizability of MPRMs for “activity cliff” and OOD molecules



Recall the two overfitting scenarios for DL-based MPRMs mentioned in the introduction: 1) struggling with ID “activity cliffs” (AC) molecules^{36,37} and 2) struggling with OOD molecules.^{35,38,53} OOD problems can be further divided into structure-OOD and property-OOD (Fig. 4a). Given MolRuleLoss’s demonstrated capacity to boost MPRM prediction accuracy, we next examined whether MolRuleLoss improves the generalizability for MPRMs when predicting AC and OOD molecules. We used a lipophilicity prediction task to assess the AC problem. Biologically, consider that an apparently simple substructure substitutions can lead to dramatic changes in lipophilicity (Fig. 4b), thus potentially preventing a compound from passing a membrane. To construct an AC test set, we selected molecule pairs that have a similarity (Morgan fingerprint) higher than 0.75 and a lipophilicity change greater than 1.0 ($\Delta\log P=1.0$, an order of magnitude difference). We then put one molecule from each pair into the training set and the other into the test set; the final test set comprised 290 molecules. The RMSE values of the GEM model with and without MolRuleLoss were 0.627 ± 0.018 and 0.674 ± 0.015 (Fig. 4c). These results demonstrate that MolRuleLoss improves the generalizability of an MPRM for the AC scenario.

Existing MPRMs frequently exhibit overfitting on OOD test sets for properties like permeability and binding affinity.^{53,54} Building on MolRuleLoss’s demonstrated performance boost based on extracted SSRs, we investigated whether the SSRs learned and selected for ID training data are generalizable for OOD prediction tasks. We used the aforementioned MP prediction task, specifically using molecules with an MP lower than 231 °C (representing the top 10% highest MP molecules) for training and using molecules with an MP greater than 231 °C for the test set (Fig. 4d). To minimize the influence of SSR errors, we specifically investigated whether incorporating SSRs with low STD values ($STD_{\max}=0.27$) via MolRuleLoss enhances MPRM accuracy. The RMSE values for the GEM model with and without MolRuleLoss on the OOD test set were 93.189 and 98.959 (Fig. 4e). Note that although the model with MolRuleLoss did outperform the model without it, the performances of both were catastrophically bad for the OOD molecules. These results support that low- STD_{\max} SSRs derived from a narrow training dataset enable an MPRM to “extrapolate” to a broader data scope.



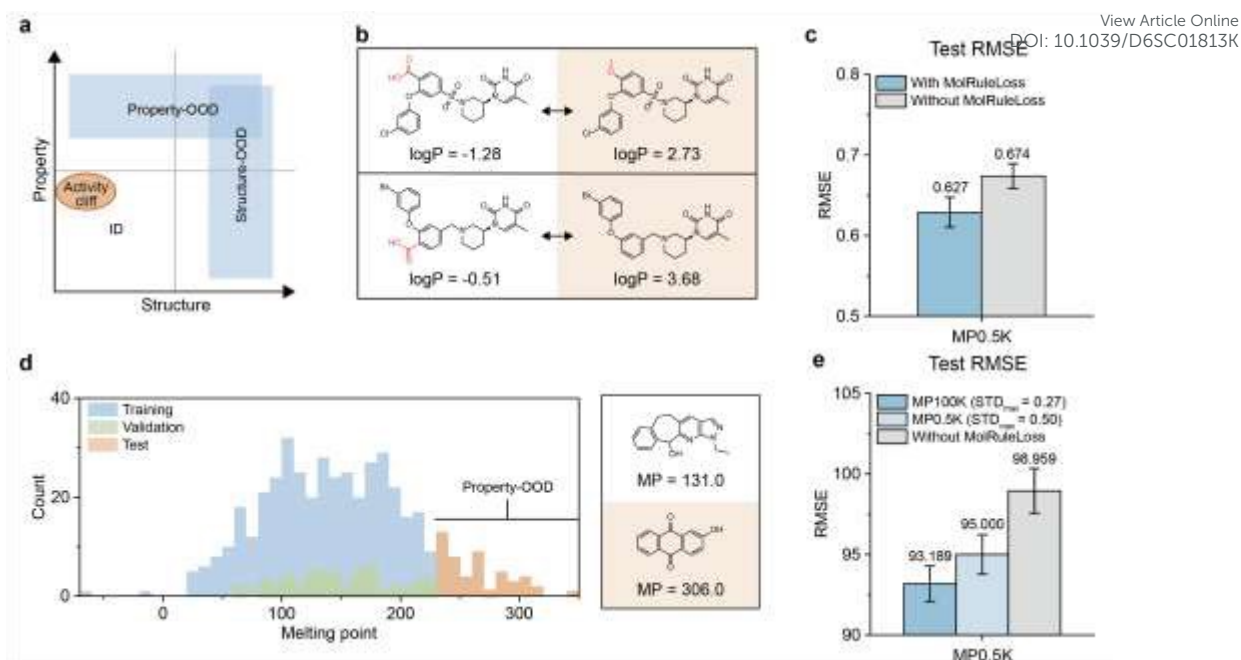


Figure 4 | Low standard deviation rules improve accuracy and generalizability of MPRMs for activity cliff and out-of-distribution molecules. **a**, Definition of “Activity cliff” (AC) and out-of-distribution (OOD) molecules. We used two axes to describe the regions for AC and OOD molecules. AC molecules are ID molecules with large property variations stemming from apparently small structural changes. Two categories of OOD molecules are considered: i. Exhibiting large structural difference to the molecules in the training dataset (“structure-OOD”), *e.g.*, molecules have no common substructures to molecules in the training dataset; ii. property values out of the range of the training dataset (“property-OOD”). **b**, Examples for the AC problem using the lipophilicity prediction task. When we use scaffold split, we already considered structure-OOD problem. To construct an AC test set, we selected molecule pairs that have a similarity (Morgan fingerprint) higher than 0.75, yet have a lipophilicity change greater than 1.0 ($\Delta\log P=1.0$, an order of magnitude difference). We then put one molecule from each pair into the training set and the other into the test set. **c**, RMSE comparison for a GEM model with or without MolRuleLoss, using the AC test set of the lipophilicity prediction task. **d**, Examples for the OOD problem using the MP prediction task. **e**, RMSE comparison for the GEM model with high-quality rules ($STD_{max}=0.27$), low-quality rules ($STD_{max}=0.50$) or without MolRuleLoss on the OOD test set.

As our low- STD_{max} rule set is neither complete nor deterministic ($STD_{max}=0$), we sought to design a task in which the rules are complete and deterministic, to quantitatively explore potential error reduction. We here employed a molecular weight



(MW) prediction, which is calculated deterministically from a molecule's SMILES string by summing atomic masses, making it an inherently error-free molecular property. MW prediction is distinct from the aforementioned lipophilicity, solubility and MP properties, which are determined using experimental measurements, and thus inevitably contain errors. We constructed an MW dataset comprising 2,700 molecules (2,160 for training and validation) by sampling 5 molecules per integer MW from 160 to 700 Dalton⁵⁵ (an MW range for drug-like molecules) from the ChEMBL dataset, with entries ranging from 160–600 Dalton as the training dataset, but ranging from 600 Dalton to 700 Dalton for the test set (*i.e.*, an OOD test set).

The RMSE value of the GEM model without MolRuleLoss was 29.507 ± 3.254 , which is catastrophically high. The 66 SSRs derived from our MW dataset represents the difference between the atomic masses of elements (delta is an exact value, $STD_{\max} = 0.000$), which are deterministic. We found that the RMSE of the GEM model with MolRuleLoss dramatically decreased, from 29.507 ± 3.254 to 0.007 ± 0.001 , and the R^2 improved from -0.050 to 1.000 (Fig. 5a,b). We also trained GEM models with and without MolRuleLoss at different data ratios, showing that much less data are required when using MolRuleLoss than without it (Fig. 5c). Additionally, when artificial noise is added to the atomic mass, which increases STD_{\max} , we observe that the RMSE values for the model with MolRuleLoss increase nearly linearly with the noise, suggesting that STD_{\max} defines the upper performance limit of our model (Fig. 5d). These results based on GEM models support that the challenges AI models currently face in handling OOD data can be addressed by acquiring SSRs from ID data and leveraging them for extrapolation to OOD contexts.



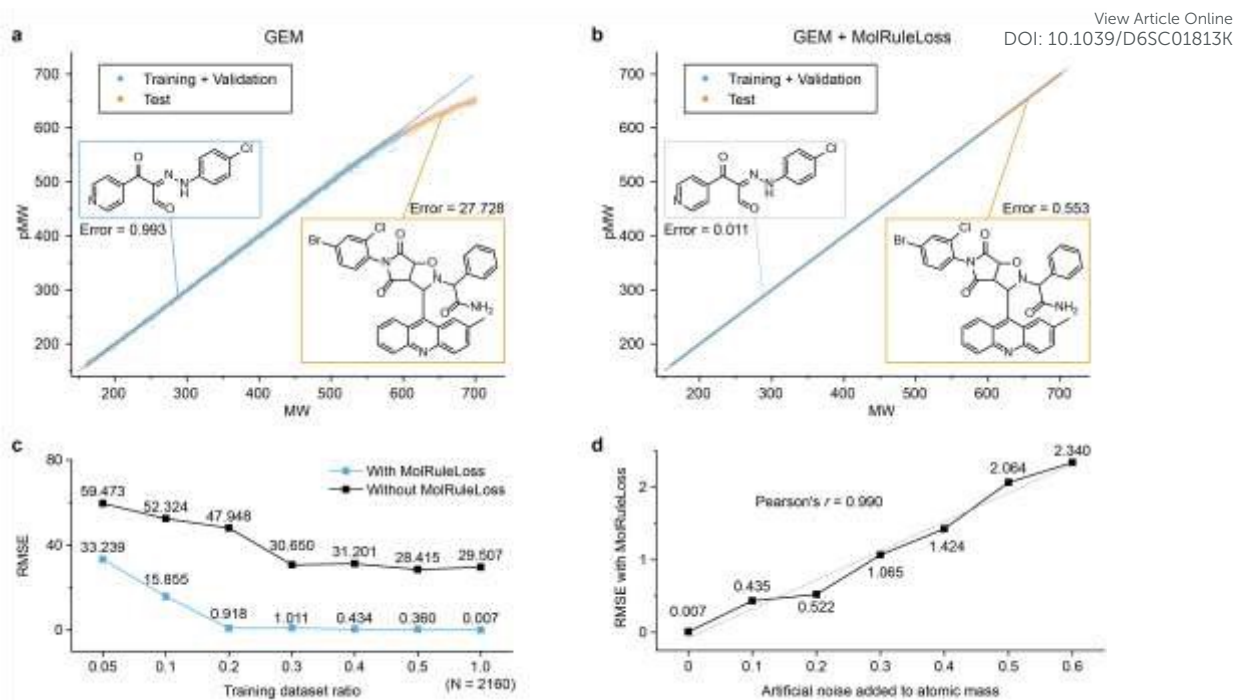


Figure 5 | Including deterministic rules with MolRuleLoss boosts the accuracy of MPRMs for the molecular weight prediction task. a-b, Molecular weight (MW) OOD prediction using the GEM model with and without MolRuleLoss. MW is calculated deterministically from a molecule's SMILES string by summing atomic masses, making it an inherently error-free property. We constructed a MW dataset comprising 2,700 molecules by sampling 5 molecules per integer MW from 160 to 700 Dalton from the ChEMBL dataset, with entries ranging from 160-600 Dalton as the training and validation dataset (in blue) but ranging from 600 Dalton to 700 Dalton for the test set (in orange). The RMSE value of the GEM model without MolRuleLoss is 29.507 ± 3.254 , which is catastrophically high; while the RMSE of the GEM model with MolRuleLoss decreased dramatically, from 29.507 ± 3.254 to 0.007 ± 0.001 . c, Performance of GEM models with and without MolRuleLoss at different data ratios, d, Performance of the GEM model with MolRuleLoss under different levels of artificial noise added to the atomic mass.

The relationship between the STD_{\max} of SSRs and the error of an MPRM

We have shown that incorporating low STD_{\max} (high-quality) SSRs via MolRuleLoss enhances the accuracy and generalizability of MPRMs. Specifically, our MP prediction task supports that MPRM with low STD_{\max} SSRs are more accurate than MPRM with



high STD_{max} SSRs; and our MW prediction task shows that at an extreme condition of $STD_{max}=0.000$, the error of the MPRM is close to 0. The large gain for the OOD with $STD_{max}=0.000$ prompted us to revisit the STD_{max} ranges for earlier prediction tasks, and we noted a trend: model error increases when STD_{max} values increase. We therefore propose our rule-variation error bound observation that the STD_{max} values of the SSRs are positively correlated with the MPRM's error. A mathematical proof supporting our rule-variation error bound observation is provided in Text S1.

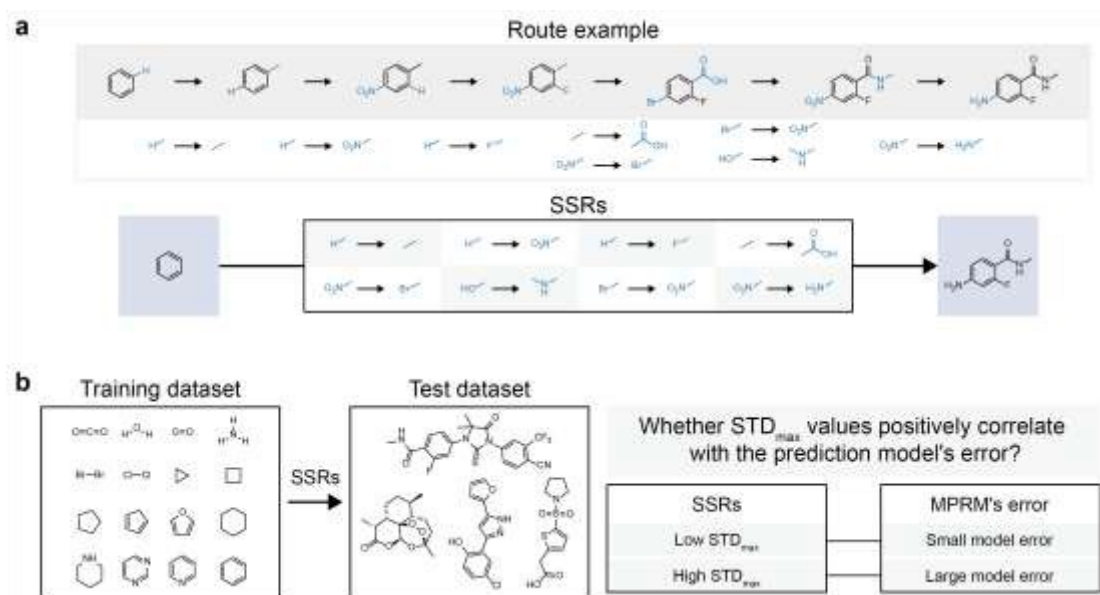


Figure 6 | Rule-variation error bound observation regarding the relationship between the STD_{max} value of SSRs and MPRM error. **a**, A given test set can be obtained by applying a series of SSRs to the training dataset. **b**, When SSRs have a large STD_{max} value, predicting a given property using an MPRM is difficult. If this property is “predictable” (that is, when the model error is smaller than a threshold value), then STD_{max} is positively correlated with this model error. We offer a mathematical proof for our rule-variation error bound observation (SI), providing a theoretical foundation for our MolRuleLoss framework.

In theory, a test set of synthesizable organic compounds could be obtained via applying SSRs to the molecules from the training dataset (Fig. 6). Our rule-variation error bound observation establishes a theoretical foundation for the MolRuleLoss framework, formally justifying the relevance of SSRs and offering guidance for future efforts to leverage such rules to improve the performance of MPRMs. Briefly, this trend emphasizes the positive utility of adding SSRs under bounded-error conditions;



moreover, there is a direct, consistent link between rule variation (STD_{max}) and model performance, offering a strong justification for using SSRs to improve model prediction accuracy. Specifically, a high STD_{max} value for SSRs of a given property indicates that an MPRM prediction for this property would have large errors⁵⁶ (e.g., as observed with binding affinity).

Discussion

In this study, we developed MolRuleLoss, a substructure-substitution-rule-informed framework that implements partial derivative constraints of SSRs derived from the MMPA method. MolRuleLoss consistently boosts the accuracy of multiple MPRMs across diverse molecular property regression tasks. It bears emphasis that even modest accuracy gains will manifest as large improvements for downstream applications, particularly in addressing data scarcity in AI-aided drug discovery. We evaluated the accuracy of MPRMs with or without MolRuleLoss across diverse molecular property regression tasks on the MoleculeNet benchmark dataset, and found that it improved the performance of SOTA MPRMs like GEM and UniMol.

In a melting point prediction task, we show that both the dimensionality and quality of rules each contribute to the magnitude of prediction accuracy gains obtained upon adding MolRuleLoss to an MPRM. Our results support that MolRuleLoss improves the generalizability of an MPRM for “activity cliff” and OOD molecules. In an OOD evaluation for MW prediction, we observed a dramatic reduction in RMSE, suggesting that MPRMs with MolRuleLoss can achieve accurate predictions when deterministic rules are included. We find that the STD_{max} values of the SSRs are positively correlated with the MPRM’s error, and we provide a mathematical proof for this rule-variation error bound observation (Text S1). Under the assumptions of model accuracy $|P(c) - F(c)| \leq e$ and model stability $|F(r(c)) - F(c)| \leq e$, the triangle inequality gives $|P(r(c)) - P(c)| \leq 3e$. The constant 3 is tight in the worst case but likely conservative in practice; distributional refinements are left for future investigation. Ultimately,

View Article Online
DOI: 10.1039/D6SC01813K



MolRuleLoss offers a framework that can be applied as a bolt-on to diverse MPRMs, supporting a wide range of applications including data scarcity scenarios in ADMET optimizations.

In principle, the MolRuleLoss approach for boosting prediction performance should be applicable for other molecular properties, such as permeability and binding affinity. Domain-knowledge integration methods operate at distinct stages of molecular learning pipelines. MolRuleLoss complements data augmentation, multitask learning, and contrastive-pretraining approaches by acting directly at the supervised objective level, and further improves activity-cliff prediction when combined with MolCLR and MolMCL (Table S3). While advances such as AlphaFold3⁵⁷ now enable prediction of protein-ligand complex structures, drug discovery efforts still require reliable ranking of candidate small molecules by binding affinity.^{58,59} Data augmentation strategies, *e.g.*, increasing data volume and adding synthetic data, have been shown to offer improvements in accuracy for MPRMs.^{60,61} However, data augmentation strategies do not necessarily improve the generalizability involving property-OOD molecules.^{35,37-39} Rule constraints seem likely to improve the generalizability of MPRMs for property predictions of OOD molecules in permeability and binding affinity tasks. Beyond GEM and UniMol, on the lipophilicity activity-cliff benchmark, integrating MolRuleLoss into MolCLR (GIN backbone) and MolMCL (GPS backbone) further reduces the 5-seed test RMSE in both cases (MolCLR: $0.649 \pm 0.032 \rightarrow 0.629 \pm 0.041$; MolMCL: $0.664 \pm 0.016 \rightarrow 0.645 \pm 0.033$), demonstrating that substructure-aware contrastive pretraining and SSR-informed loss constraints provide complementary, rather than redundant, inductive biases (Table S4). We anticipate that MolRuleLoss will likely improve generalizability when added onto other MPRMs such as VisNet and attentive FP.^{20,62}

Our work highlights several directions for enhancing the accuracy of MPRMs through rule-based partial derivative constraints. We have so far focused on SSRs derived from the MMPA method, and the number of rules is limited when the training dataset size is small. The effectiveness of MolRuleLoss depends on the availability and reliability of SSRs. On the melting-point task, transferring 175 SSRs from MP_100K



to MP_0.5K improved RMSE by 8.6%, compared with 4.4% using only the 44 SSRs extracted in-domain. Future work may explore the generation or learning of additional SSR-like constraints from larger datasets or model-assisted substructure searches. Future extensions could incorporate alternative rule sources. Notably, while we have currently focused solely on pairwise substructure substitution, the MolRuleLoss framework should accommodate other rule formats, for example those involving substitution of more than one substructure. Beyond the domain of molecular property prediction, MolRuleLoss may be applied to other regression tasks (*e.g.*, oral bioavailability). Finally, MolRuleLoss may be applicable to molecular entities including proteins or nucleic acids. To achieve this expanded application scope, one could define substructures of proteins and use AI models to learn potential SSRs from a supervised model.

Methods

Dataset

MoleculeNet

We sourced molecular datasets from MoleculeNet,⁵⁰ focusing on three benchmark tasks: Lipophilicity (logP, n=4,200), ESOL solubility (logS, n=1,150), and FreeSolv hydration free energies (ΔG_{solv} , n=650). Molecular structures were standardized using RDKit (v2022.09)⁶³ through a multi-step pipeline that included charge neutralization, salt removal, and aromaticity normalization to Kekulé form. To rigorously evaluate out-of-distribution (OOD) generalization, we implemented two complementary splitting strategies. For scaffold-based OOD evaluation, molecules were clustered by Bemis-Murcko scaffolds using Butina clustering with a Tanimoto similarity cutoff of 0.7 based on Morgan fingerprints (radius 2), with the 10% least populated clusters reserved for testing to maximize structural novelty. For property-based OOD evaluation, molecules were sorted by target property values, with the top and bottom 10% extremes forming the test set to simulate prediction challenges for property outliers. To specifically test extreme extrapolation capabilities, we constructed a dedicated molecular weight

View Article Online
DOI: 10.1039/D6SC01813K



prediction dataset from ChEMBL v30,⁶⁴ spanning 160–700 Da with uniform sampling, where the test set exclusively contained molecules >600 Da (n=300) and inputs were restricted to counts of 12 common atoms (H, C, N, O, F, P, S, Cl, Br, I, B, Si).

Melting Point Dataset

Our melting point (MP) dataset was extracted from USPTO patents (1976-2014) using Tetko et al.'s text-mining pipeline.⁶⁵ Suspicious records were excluded, including MPs >500 °C, ranges >50 °C, or inverted ranges, removing 1,498 entries from grants and 426 from applications. Duplicate entries ($\Delta\text{MP} \leq 1$ °C) were consolidated, eliminating 366,532 records. For compounds with multiple valid measurements, the value closest to the median was retained to preserve data integrity and patent traceability. Molecules failing descriptor calculation were excluded. Experimental accuracy was estimated at $\sigma = 35$ °C based on 18,058 duplicate measurements, accounting for polymorphism and extraction variances. For the drug-like space (50–250 °C, covering >90% of compounds), error refined to $\sigma = 32$ °C. The final curated dataset provides a diverse MP collection for predictive modeling. Overlap with established benchmarks was removed to ensure test set independence.

Butina Clustering Implementation

We applied the Butina clustering algorithm to identify structurally homogeneous groups within our molecular dataset.⁶⁶ This method partitions compounds based on pairwise structural similarity, ensuring that all molecules within a cluster exhibit a minimum similarity to the cluster centroid. Compounds in SMILES format were converted into binary RDKit RDKit5 fingerprint vectors, capturing key substructural features. Pairwise structural similarities were computed using the Tanimoto coefficient (T_c):

$$T_c = \frac{A \cap B}{A \cup B} \quad (\text{Eqn.2})$$

This generated a symmetric similarity matrix, subsequently converted to a distance matrix via $D = 1 - T_c$.

Molecules were ranked by their number of neighbors (compounds within a



similarity threshold). The molecule with the most unassigned neighbors was selected as a centroid, and all its neighbors were assigned to the new cluster. These members were then excluded from future centroid selection or cluster membership. Remaining unassigned molecules were classified as singletons. This algorithm efficiently produces clusters with high internal structural consistency.

View Article Online
DOI: 10.1039/D6SC01813K

Model Architectures

To enhance the generalization capability of molecular property prediction, we integrated Physics-Informed Neural Networks (PINNs)⁴⁶ into the molecular representation learning framework. PINNs represent a class of neural networks that incorporate prior physical knowledge, typically expressed as partial differential equations (PDEs), ordinary differential equations (ODEs), or other physical constraints, directly into the learning process through a multi-component loss function. This approach reduces the reliance on large volumes of purely experimental data and ensures that predictions adhere to fundamental physical laws. In our implementation, building upon established molecular representation learning models, we designed a feedforward neural network (FNN) architecture to serve as a universal function approximator mapping molecular substructure counts to target properties (*e.g.*, logP, logS, ΔG_{solv}). The core innovation of PINNs lies in their composite loss function, which simultaneously optimizes both data fidelity and physical consistency.

Two state-of-the-art architectures served as baselines for molecular representation learning. The Geometry-Enhanced Molecular GNN (GEM)¹⁷ processes 2D molecular graphs through an 8-layer Graph Isomorphism Network (GIN) with 32-dimensional hidden states, using atom features that encode atomic number, degree, hybridization, implicit valence, and formal charge. Pretrained via self-supervised bond length and angle prediction tasks, GEM employs a two-layer MLP (32 \rightarrow 32 \rightarrow 1) with dropout ($p=0.1$) for property regression. In contrast, UniMol (Universal Molecular Transformer)¹⁹ operates on 3D conformers generated using RDKit's ETKDG method with MMFF94 optimization. This SE(3)-equivariant Transformer architecture uses



512-dimensional embeddings with 8 attention heads and rotary positional embeddings, pretrained on 209 million molecular conformations from QM9 and GEOM datasets. For property prediction, UniMol reduces dimensions through a $512 \rightarrow 128 \rightarrow 1$ MLP with LayerNorm. Both models were implemented in PyTorch with mixed-precision training on NVIDIA A100 GPUs.

MolRuleLoss Framework

Our approach integrates chemical domain knowledge through differentiable substructure substitution rules (SSRs), systematically derived from matched molecular pair (MMP) analysis. Using RDKit's MMP implementation (max cut size=13 atoms), we identified all molecular pairs differing by exactly one substructure and computed mean property changes (ΔP) and their standard deviations (σ) for each substitution pattern. Twenty-nine high-confidence rules were retained ($\sigma < 0.5$ and occurrence ≥ 10), including common transformations like phenyl \rightarrow cyclohexyl ($\Delta \log P = -0.82 \pm 0.21$) and $-\text{OH} \rightarrow -\text{NH}_2$ ($\Delta \log S = +0.47 \pm 0.33$). The MolRuleLoss function combines conventional prediction error with a rule consistency term.

For a rule $r = (x_i \rightarrow x_j)$, let Δ_r denote the mean property shift obtained from matched molecular pairs, and let h_m represent the molecular graph embedding of molecule m . The adaptive rule term is defined as

$$\Delta_r^{\text{adaptive}}(m) = \Delta_r + h_m^T \theta_{x_i \rightarrow x_j} \quad (\text{Eqn.3})$$

where $\theta \in \mathbb{R}^{\{|V| \times d\}}$ is a shared learnable parameter matrix indexed over the substructure vocabulary ($|V|$ is the substructure vocabulary size, d is the embedding dimension). The rule residual for molecule m is then expressed as

$$\varepsilon_r(m) = \left(\frac{\partial \hat{y}_m}{\partial \#x_i} - \frac{\partial \hat{y}_m}{\partial \#x_j} \right) - \Delta_r^{\text{adaptive}}(m) \quad (\text{Eqn.4})$$

The SSR objective is computed as

$$\text{Loss}_{SSR} = \frac{1}{|R|} \sum_{r \in R} \varepsilon_r(m)^2 \quad (\text{Eqn.5})$$

and the final optimization objective is given by”

$$\text{Loss} = \text{Loss}_{MSE} + \lambda \text{Loss}_{SSR} \quad (\text{Eqn.1})$$



Recognizing that the impact of identical substructural substitutions on molecular properties can be significantly modulated by their distinct chemical environments, we introduce an adaptive rule refinement algorithm. This approach addresses the critical limitation of conventional matched molecular pair (MMP) analysis, wherein static transformation rules fail to account for context-dependent effects on property changes.

Our methodology dynamically adjusts rule-based predictions by integrating learned molecular representations with the specific contextual environment of each molecule. The implementation proceeds as follows: For a given substructural transformation rule with an initial predicted property change value Δ , we generate a high-dimensional molecular representation vector h using a pre-trained molecular graph neural network (GNN). This representation comprehensively encodes the structural and electronic features of the entire molecular scaffold surrounding the substitution site.

Simultaneously, we initialize a learnable parameter vector θ of identical dimension to the molecular representation. The adaptive adjustment term is computed as the dot product between the molecular representation vector and the parameter vector:

$$\delta = h \cdot \theta \quad (\text{Eqn.6})$$

This term is then added to the base rule value to yield the context-sensitive prediction:

$$\Delta_{\text{Adaptive}} = \Delta + \delta \quad (\text{Eqn.7})$$

During the model training phase, both the molecular representation network (if fine-tuned) and the adaptive parameter vectors are optimized jointly using backpropagation to minimize the difference between predicted and true property changes. This learning process enables the model to capture systematic variations in rule applicability across different chemical environments.

Upon convergence, the learned parameter vectors become fixed and are utilized in all subsequent predictions. This approach effectively transforms static, context-agnostic rules into dynamic, adaptive rules that respond to the specific chemical environment of



each molecular transformation, thereby significantly enhancing prediction accuracy for complex molecular systems.

Training and Evaluation

All models were trained using Adam optimization ($\text{lr}=1\text{e-}4$, $\beta_1=0.900$, $\beta_2=0.999$) with decoupled weight decay ($1\text{e-}5$) and cosine annealing with warm restarts (15 epochs per cycle). Training employed batch size 32 with dropout ($p=0.1$), gradient clipping (max norm=5.0), and early stopping (10-epoch patience). Performance was assessed via root mean squared error (RMSE) and coefficient of determination (R^2), with five replicates (seeds 1024–1028) to estimate variance. Specialized OOD tests included activity cliff evaluation (290 molecular pairs with Tanimoto similarity >0.75 but $\Delta\log P >1.0$) and molecular weight extrapolation (train: 160–600 Dalton, test: 600–700 Dalton). MolRuleLoss-specific hyperparameters (λ and STD_max) were optimized using a standard Optuna TPE search on the train/validation split while holding out the test set for final evaluation; the released script `src/search_lipo_gin.py` provides a reusable automated search pipeline for new datasets. Full reproducibility is ensured through publicly available code, datasets, and configuration files that document all preprocessing steps, model implementations, and hyperparameter choices.

GEM Model Architecture

The core of our model utilizes a multi-layer GNN architecture based on the message-passing framework. At each layer l , the node representations are updated through the following operations:

Message Passing: For node $v \in V$ at layer l , messages from neighbors $u \in N(v)$ are computed as:

$$m_{u \rightarrow v}^{(l)} = \text{MSG}^{(l)}(h_u^{(l-1)}, h_v^{(l-1)}, e_{uv}) \quad (\text{Eqn.8})$$

where $N(v)$ denotes the neighborhood of node v , $h_u^{(l-1)}$ and $h_v^{(l-1)}$ are node features from the previous layer, and e_{uv} represents edge features.



Aggregation: The aggregated message for node v is obtained by:

View Article Online
DOI: 10.1039/D6SC01813K

$$M_v^{(l)} = \text{AGGREGATE}^{(l)}(\{m_{u \rightarrow v}^{(l)} : u \in N(v)\}) \quad (\text{Eqn. 9})$$

We employ sum aggregation followed by a multilayer perceptron (MLP), similar to the Graph Isomorphism Network (GIN) approach:

$$a_v^{(l)} = \sum_{u \in N(v)} (h_u^{(l-1)} + h_v^{(l-1)} + e_{uv}) \quad (\text{Eqn. 10})$$

Update: The node representation is updated through a COMBINE function:

$$h_v^{(l)} = \text{COMBINE}^{(l)}(h_v^{(l-1)}, M_v^{(l)}) = \text{MLP}^{(l)}(a_v^{(l)}) \quad (\text{Eqn. 11})$$

where $\text{MLP}^{(l)}$ denotes a 2-layer multilayer perceptron with hidden size of 32.

The model consists of 8 GNN layers with hidden dimension 32, providing sufficient depth for capturing complex molecular patterns. Layer Normalization applied before each message passing step. Graph Size Normalization to address variable graph sizes. Residual Connections to facilitate gradient flow in deep architectures.

For graph-level prediction tasks, we generate a global graph representation through a readout function that aggregates all node representations from the final GNN layer:

$$h_G = \text{READOUT}(\{h_v^{(l)} : v \in V\}) \quad (\text{Eqn. 12})$$

We employ global mean pooling as the readout function, resulting in a 32-dimensional graph embedding.

To enhance predictive performance, we integrate domain-specific molecular substructure features with the learned graph representation. Let $f_{\text{substructure}}$ denote the m -dimensional vector encoding the counts of various molecular substructures. The final integrated representation is obtained by concatenation:

$$h_{\text{final}} = \text{Concat}(h_G, f_{\text{substructure}}) \quad (\text{Eqn. 13})$$

This combined representation is then processed through a prediction head consisting of



a 3-layer MLP with hidden size of 128:

$$\hat{y} = \text{MLP}_{\text{prediction}}(h_{\text{final}}) \quad (\text{Eqn. 14})$$

where the output dimension depends on the specific prediction task. This architecture allows the model to simultaneously leverage both learned graph embeddings and expert-defined molecular features for enhanced predictive performance.

Matched Molecular Pair Analysis (MMPA) Implementation

We employed the mmpdb package (v3.1.3) for Matched Molecular Pair Analysis⁶⁷ to systematically extract structural transformation rules and associated property change profiles from our molecular dataset. The workflow consisted of four phases:

Compounds in SMILES format were processed using the mmpdb fragmentation algorithm, which cuts 1-3 non-ring bonds to decompose structures into a constant core fragment and variable R-group fragments. This captures all possible matched molecular pairs (MMPs). Fragmented representations were indexed to identify compound pairs sharing identical constant cores but differing in one variable fragment. This generated a database of structural transformations, with each transformation cataloged alongside its experimental property measurements. Each transformation was characterized by its rule environment using two representations: 1) SMARTS patterns (RDKit circular fingerprints, radius 2) for machine-readable environment definitions; 2) Pseudo-SMILES strings for human-interpretable representations.

This enables differentiation of transformations across distinct chemical contexts. For each transformation-environment combination, we computed mean property change, standard deviation, standard error, and fraction of pairs exceeding property change thresholds. Transformations with ≥ 10 matched pairs and well-defined distributions (low standard error) were retained as predictive rules.

Author Contributions

Xiaoyu Fan: Methodology, Software, Investigation, Validation, Writing – original draft;



Lin Guo: Visualization; Ruizhen Jia: Writing – review & editing; Yang Tian: Conceptualization, Methodology; Zhihao Yang: Formal analysis, Data curation; Weihao Li: Formal analysis, Validation; Boxue Tian: Conceptualization, Methodology, Supervision, Writing – original draft, Writing – review & editing, Project administration, Funding acquisition.

View Article Online
DOI: 10.1039/D6SC01813K

Conflicts of interests

The authors declare no competing interests.

Data Availability

The datasets of molecular properties affected by this work can be found in the public website MoleculeNet at <https://moleculenet.org/>. MW dataset can be found in the Github repository at <https://github.com/fanxiaoyu0/MolRuleLoss>.

Code Availability

The code used for this paper can be found in the Github repository at <https://github.com/fanxiaoyu0/MolRuleLoss>.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.22473067), Beijing Frontier Research Center for Biological Structure (No. 041500002), Tsinghua University Initiative Scientific Research Program (No.20231080030), and the Tsinghua-Peking University Center for Life Sciences (No.20111770319).

Author Information

Correspondence and requests for materials should be addressed to B.T. (boxuetian@mail.tsinghua.edu.cn).

Notes and References



- [1] Z. Niu, X. Xiao, W. Wu, Q. Cai, Y. Jiang, W. Jin, M. Wang, G. Yang, L. Kong, X. Jin, G. Yang and H. Chen, *Sci. Data*, 2024, **11**, 985.
- [2] K. Swanson, P. Walther, J. Leitz, S. Mukherjee, J. C. Wu, R. V. Shivnaraine and J. Zou, *Bioinformatics*, 2024, **40**, btae416.
- [3] J. Yi, S. Shi, L. Fu, Z. Yang, P. Nie, A. Lu, C. Wu, Y. Deng, C. Hsieh, X. Zeng, T. Hou and D. Cao, *Nat. Protoc.*, 2024, **19**, 1105-1121.
- [4] L. Fu, S. Shi, J. Yi, N. Wang, Y. He, Z. Wu, J. Peng, Y. Deng, W. Wang, C. Wu, A. Lyu, X. Zeng, W. Zhao, T. Hou and D. Cao, *Nucleic Acids Res.*, 2024, **52**, W422-W431.
- [5] Y. Gu, Z. Yu, Y. Wang, L. Chen, C. Lou, C. Yang, W. Li, G. Liu and Y. Tang, *Nucleic Acids Res.*, 2024, **52**, W432-W438.
- [6] W. P. Walters and R. Barzilay, *Acc. Chem. Res.*, 2021, **54**, 263-270.
- [7] X. Zeng, H. Xiang, L. Yu, J. Wang, K. Li, R. Nussinov and F. Cheng, *Nat. Mach. Intell.*, 2022, **4**, 1004-1016.
- [8] Y. Zheng, H. Y. Koh, J. Ju, A. T. Nguyen, L. T. May, G. I. Webb and S. Pan, *Nat. Mach. Intell.*, 2025, **7**, 437-447.
- [9] H. Shimakawa, A. Kumada and M. Sato, *npj Comput. Mater.*, 2024, **10**, 11.
- [10] E. Nuñez-Andrade, I. Vidal-Daza, R. Gomez-Bombarelli, J. W. Ryan and F. J. Martin-Martinez, *ChemRxiv*, 2025, Preprint, DOI: 10.26434/chemrxiv-2025-6hfp8.
- [11] M. Gallegos, V. Vassilev-Galindo, I. Poltavsky, Á. Martín Pendás and A. Tkatchenko, *Nat. Commun.*, 2024, **15**, 4345.
- [12] A. Rasool, J. Ul Rahman and R. Uwitije, *J. Cheminform.*, 2025, **17**, 81.
- [13] Z. Wu, J. Wang, H. Du, D. Jiang, Y. Kang, D. Li, P. Pan, Y. Deng, D. Cao, C.-Y. Hsieh and T. Hou, *Nat. Commun.*, 2023, **14**, 2585.
- [14] D. Chen, K. Gao, D. D. Nguyen, X. Chen, Y. Jiang, G.-W. Wei and F. Pan, *Nat. Commun.*, 2021, **12**, 3521.
- [15] J. Qiao, J. Jin, D. Wang, S. Teng, J. Zhang, X. Yang, Y. Liu, Y. Wang, L. Cui, Q. Zou, R. Su and L. Wei, *Nat. Commun.*, 2025, **16**, 4382.
- [16] S. Wang, Y. Guo, Y. Wang, H. Sun and J. Huang, *In Proc 10th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics*, 2019.
- [17] X. Fang, L. Liu, J. Lei, D. He, S. Zhang, J. Zhou, F. Wang, H. Wu and H. Wang, *Nat. Mach. Intell.*, 2022, **4**, 127-134.
- [18] Z. Yang, L. Wang, T. Huang, Y. Wang, M. Gao, T. Hou, J. Ding and J. Xiao, *Adv. Sci.*, 2025, 1-11.
- [19] G. Zhou, Z. Gao, Q. Ding, H. Zheng, H. Xu, Z. Wei, L. Zhang and G. Ke, *In Proc. 11th International Conference on Learning Representations*, 2023.
- [20] Y. Wang, T. Wang, S. Li, X. He, M. Li, Z. Wang, N. Zheng, B. Shao and T.-Y. Liu, *Nat. Commun.*, 2024, **15**, 313.
- [21] L. Liu, D. He, X. Fang, S. Zhang, F. Wang, J. He and H. Wu, *arXiv*, 2022, Preprint, DOI: arxiv-2208.05863.
- [22] X. Ji, Z. Wang, Z. Gao, H. Zheng, L. Zhang, G. Ke and W. E, *In 38th Conference on Neural Information Processing Systems*, 2024.
- [23] B. Dou, Z. Zhu, E. Merkurjev, L. Ke, L. Chen, J. Jiang, Y. Zhu, J. Liu, B. Zhang and G.-W. Wei, *Chem. Rev.*, 2023, **123**, 8736-8780.



- [24] T. Yin, G. Panapitiya, E. D. Coda and E. G. Saldanha, *J. Cheminform.*, 2023, **15**, 105. [View Article Online](#)
DOI: 10.1039/D6SC01813K
- [25] Y. Ji, L. Zhang, J. Wu, B. Wu, L. Li, L.-K. Huang, T. Xu, Y. Rong, J. Ren, D. Xue, H. Lai, W. Liu, J. Huang, S. Zhou, P. Lou, P. Zhao and Y. Bian *In Proc 37th AAAI Conference on Artificial Intelligence*, 2023.
- [26] J. Kim, J. Willette, B. Andreis and S. J. Hwang, *arXiv*, 2025, Preprint, DOI: arxiv-2506.11877.
- [27] W. X. Shen, C. Cui, X. Su, Z. Zhang, A. Velez-Arce, J. Wang, X. Shi, Y. Zhang, J. Wu, Y. Z. Chen and M. Zitnik, *ChemRxiv*, 2024, Preprint, DOI: 10.26434/chemrxiv-2023-5cz7s-v2.
- [28] M. Dablander, *ArXiv*, 2024, Preprint, DOI: arxiv-2411.13688.
- [29] Y. Zhang, X. Li, N. Peng, Y. Chen, H. Xia and Y. Gu, *In Proc 31st International Conference on Neural Information Processing*, 2024, pp. 31-48.
- [30] Y. Wan, J. Wu, T. Hou, C.-Y. Hsieh and X. Jia, *Nat. Commun.*, 2025, **16**, 413.
- [31] J. Jiang, R. Zhang, Y. Yuan, T. Li, G. Li, Z. Zhao and Z. Yu, *J Mol. Graph. Model.*, 2023, **121**, 108454.
- [32] Z. Wang, T. Jiang, J. Wang, J. Shao, B. Wei and Q. Xuan and H. Wang, *IEEE Transactions on Computational Biology and Bioinformatics*, 2025.
- [33] Z. Cheng, H. Xiang, P. Ma, L. Zeng, X. Jin, X. Yang, J. Lin, Y. Deng, B. Song and X. Feng, Y. Deng, C. Deng, B. Song and X. Zeng, *BMC Biol.*, 2025, **23**, 279
- [34] T. Kuang, P. Liu and Z. Ren, *Big Data Min. Anal.*, 2024, **7**, 858-888.
- [35] E. R. Antoniuk, S. Zaman, T. Ben-Nun, P. Li, J. Diffenderfer, B. Demirci, O. Smolenski, T. Hsu, A. M. Hiszpanski, K. Chiu, B. Kailkhura and B. Van Essen, *arXiv*, 2025, Preprint, DOI: arxiv-2505.01912.
- [36] D. Van Tilborg, A. Alenicheva and F. Grisoni, *J. Chem. Inf. Model.*, 2022, **62**, 5938-5951.
- [37] S. Tamura, T. Miyao and J. Bajorath, *J. Cheminform.*, 2023, **15**, 4.
- [38] J. Deng, Z. Yang, H. Wang, I. Ojima, D. Samaras and F. Wang, *Nat. Commun.*, 2023, **14**, 6395.
- [39] C. Pang, H. H. Tong and L. Wei, *Quant. Biol.*, 2023, **11**, 395-404.
- [40] Y. Wang, J. Wang, Z. Cao and A. Barati Farimani, *Nat. Mach. Intell.*, 2022, **4**, 279-287.
- [41] Y. Wan, J. Wu, T. Hou, C. Y. Hsieh and X. Jia, *Nat. Commun.*, 2025, **16**, 413.
- [42] Y. Fang, Q. Zhang, N. Zhang, Z. Chen, X. Zhuang, X. Shao, X. Fan and H. Chen, *Nat. Mach. Intell.*, 2023, **5**, 542-553.
- [43] J. Terven, D.-M. Cordova-Esparza, J.-A. Romero-González, A. Ramírez-Pedraza and E. Chávez-Urbiola, *Artif. Intell. Rev.*, 2025, **58**, 195.
- [44] E. Heid, K. P. Greenman, Y. Chung, S.-C. Li, D. E. Graff, F. H. Vermeire, H. Wu, W. H. Green and C. J. McGill, *J. Chem. Inf. Model.*, 2024, **64**, 9-17.
- [45] Y. Qin, C. Li, X. Shi and W. Wang, *Front. Bioeng. Biotechnol.*, 2022, **10**, 946329.
- [46] S. Wang, S. Sankaran, H. Wang and P. Perdikaris, *arXiv*, 2025, Preprint, DOI: arxiv-2308.08468.
- [47] A. Krishnapriyan, A. Gholami, S. Zhe, R. Kirby and M. W. Mahoney, *In Proc. 35th Conference on Neural Information Processing Systems*, 2021.
- [48] Z. Yang, S. Shi, L. Fu, A. Lu, T. Hou and D. Cao, *J. Med. Chem.*, 2023, **66**, 4361-4377.
- [49] K. Yang, K. Swanson, W. Jin, C. Coley, P. Eiden, H. Gao, A. Guzman-Perez, T. Hopper, B. Kelley, M. Mathea, A. Palmer, V. Settels, T. Jaakkola, K. Jensen and R. Barzilay, *J. Chem.*



- Inf. Model.*, 2019, **59**, 3370-3388.
- [50] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing and V. Pande, *Chem. Sci.*, 2018, **9**, 513-530.
- [51] I. Sushko, S. Novotarskyi, R. Körner, A. K. Pandey, M. Rupp, W. Teetz, S. Brandmaier, A. Abdelaziz, V. V. Prokopenko, V. Y. Tanchuk, R. Todeschini, A. Varnek, G. Marcou, P. Ertl, V. Potemkin, M. Grishina, J. Gasteiger, I. I. Baskin, V. A. Palyulin, E. V. Radchenko, W. J. Welsh, V. Kholodovych, D. Chekmarev, A. Cherkasov, J. Aires-de-Sousa, Q.-Y. Zhang, A. Bender, F. Nigsch, L. Patiny, A. Williams, V. Tkachenko and I. V. Tetko, *J. Comput. Aided Mol. Des.*, 2011, **25**, 533-554.
- [52] K. Weiss, T. M. Khoshgoftaar and D. Wang, *J. Big Data*, 2016, **3**, 9.
- [53] H. Fooladi, T. N. L. Vu, M. Mathea and J. Kirchmair, *J. Chem. Inf. Model.*, 2025, **65**, 19, 9871-9891
- [54] Y. Shi, W. Xu and P. Hu, *Briefings Bioinf.*, 2025, **26**, bbaf294.
- [55] G. R. Bickerton, G. V. Paolini, J. Besnard, S. Muresan and A. L. Hopkins, *Nat. Chem.*, 2012, **4**, 90-98.
- [56] H. Zhang, X. Liu, W. Cheng, T. Wang and Y. Chen, *Comput. Biol. Med.*, 2024, **174**, 108435.
- [57] J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A. J. Ballard, J. Bambrick, S. W. Bodenstein, D. A. Evans, C.-C. Hung, M. O'Neill, D. Reiman, K. Tunyasuvunakool, Z. Wu, A. Žemgulytė, E. Arvaniti, C. Beattie, O. Bertolli, A. Bridgland, A. Cherepanov, M. Congreve, A. I. Cowen-Rivers, A. Cowie, M. Figurnov, F. B. Fuchs, H. Gladman, R. Jain, Y. A. Khan, C. M. R. Low, K. Perlin, A. Potapenko, P. Savy, S. Singh, A. Stecula, A. Thillaisundaram, C. Tong, S. Yakneen, E. D. Zhong, M. Zielinski, A. Židek, V. Bapst, P. Kohli, M. Jaderberg, D. Hassabis and J. M. Jumper, *Nature*, 2024, **630**, 493-500.
- [58] Z. Chen, B. Peng, T. Zhai, D. Adu-Ampratwum and X. Ning, *Nat. Comput. Sci.*, 2024, **4**, 899-909.
- [59] L. An, M. Said, L. Tran, S. Majumder, I. Goreshnik, G. R. Lee, D. Juergens, J. Dauparas, I. Anishchenko, B. Coventry, A. K. Bera, A. Kang, P. M. Levine, V. Alvarez, A. Pillai, C. Norn, D. Feldman, D. Zorine, D. R. Hicks, X. Li, M. Garcia Sanchez, D. K. Vafeados, P. J. Salvesson, A. A. Vorobieva and D. Baker, *Science*, 2024, **385**, 276-282.
- [60] M. bin Javaid, T. Gervens, A. Mitsos, M. Grohe and J. G. Rittig, *Comput. Chem. Eng.*, 2025, **201**, 109253.
- [61] J. Jiang, Y. Li, R. Zhang and Y. Liu, *J. Mol. Graph. Model.*, 2024, **128**, 108703.
- [62] Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang and M. Zheng, *J. Med. Chem.*, 2019, **63**, 8749-8760.
- [63] G. Landrum, *Release*, 2013, **1**, 4.
- [64] A. Gaulton, L. J. Bellis, A. P. Bento, J. Chambers, M. Davies, A. Hersey, Y. Light, S. McGlinchey, D. Michalovich, B. Al-Lazikani and P. Overington, *Nucleic Acids Res.*, 2012, **40**, D1100-D1107.
- [65] I. V. Tetko, D. M. Lowe and A. J. Williams, *J. Cheminf.*, 2016, **8**, 2.
- [66] D. Butina, *J. Chem. Inf. Model.*, 1999, **39**, 747-750.
- [67] A. G. Dossetter, E. J. Griffen and A. G. Leach, *Drug Discov. Today*, 2013, **18**, 724-731.



Data Availability Statement (DAS)

View Article Online
DOI: 10.1039/D6SC01813K

Data Availability

The datasets of molecular properties affected by this work can be found in the public website MoleculeNet at <https://moleculenet.org/>. MW dataset can be found in the Github repository at <https://github.com/fanxiaoyu0/MolRuleLoss>.

Code Availability

The code used for this paper can be found in the Github repository at <https://github.com/fanxiaoyu0/MolRuleLoss>.

