Transfer learning for molecular property predictions from small datasets

Cite as: AIP Advances 14, 105119 (2024); doi: 10.1063/5.0214754
Submitted: 15 August 2024 • Accepted: 30 September 2024 •
Published Online: 14 October 2024

View Online Export Citation CrossMark

Thorren Kirschbaum[1] and Annika Bande[1,2,a)]

AFFILIATIONS

[1] Theory of Electron Dynamics and Spectroscopy, Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Hahn-Meitner-Platz 1, 14109 Berlin, Germany
[2] Institute of Inorganic Chemistry, Leibniz University Hannover, Callinstr. 9, 30167 Hannover, Germany

[a)] **Author to whom correspondence should be addressed:** annika.bande@helmholtz-berlin.de

ABSTRACT

Machine learning has emerged as a new tool in chemistry to bypass expensive experiments or quantum-chemical calculations, for example, in high-throughput screening applications. However, many machine learning studies rely on small datasets, making it difficult to efficiently implement powerful deep learning architectures such as message passing neural networks. In this study, we benchmark common machine learning models for the prediction of molecular properties on two small datasets, for which the best results are obtained with the message passing neural network PaiNN as well as SOAP molecular descriptors concatenated to a set of simple molecular descriptors tailored to gradient boosting with regression trees. To further improve the predictive capabilities of PaiNN, we present a transfer learning strategy that uses large datasets to pre-train the respective models and allows us to obtain more accurate models after fine-tuning on the original datasets. The pre-training labels are obtained from computationally cheap *ab initio* or semi-empirical models, and both datasets are normalized to mean zero and standard deviation one to align the labels' distributions. This study covers two small chemistry datasets, the Harvard Organic Photovoltaics dataset (HOPV, HOMO–LUMO-gaps), for which excellent results are obtained, and the FreeSolv dataset (solvation energies), where this method is less successful, probably due to a complex underlying learning task and the dissimilar methods used to obtain pre-training and fine-tuning labels. Finally, we find that for the HOPV dataset, the final training results do not improve monotonically with the size of the pre-training dataset, but pre-training with fewer data points can lead to more biased pre-trained models and higher accuracy after fine-tuning.

© 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/). https://doi.org/10.1063/5.0214754

06 January 2025 06:56:56

I. INTRODUCTION

In the realm of molecular sciences, predicting the properties of molecules stands as a pivotal pursuit across numerous domains, including pharmaceuticals, materials science, and environmental studies. Traditionally, this endeavor has been tackled by quantum chemistry, relying on sophisticated theoretical frameworks and computational simulations to unravel the intricate behaviors of molecules. However, the advent of machine learning (ML) has ushered in a paradigm shift, offering an alternative approach that leverages data-driven methodologies to efficiently compute molecular properties.[1–5]

One notable application of ML in this domain is its integration into high throughput screening workflows, where it replaces conventional quantum chemistry methods.[3] Through the judicious application of ML algorithms, researchers can expedite the process of identifying promising molecules for drug discovery or material development, thereby accelerating the pace of scientific innovation. However, the powerful deep neural networks (NNs) often used in such approaches usually need large amounts of training data to finally make accurate predictions.

Machine learning algorithms, particularly deep NNs, have exhibited remarkable success in discerning intricate patterns and relationships within datasets of substantial size. In recent developments, the message passing neural network PaiNN has demonstrated the highest performance in predicting molecular properties, achieving state-of-the-art results across various datasets.[6,7] Another prevalent approach in ML for molecular sciences entails the utiliza-

© Author(s) 2024

tion of tailored molecular descriptors, such as SOAP,[8,9] in conjunction with traditional ML models like kernel ridge regression (KRR). In general, as a rule of thumb, deep learning architectures are to be preferred when $>10^3$–10^4 training points are available, and different approaches, such as KRR, should be considered otherwise.[10] Accordingly, molecular fingerprinting methods are commonly used in junction with KRR for regression on small datasets.[7,11–13] Yet, in the context of molecular property prediction, datasets often fall short regarding the typically required large amounts of data for deep NNs. This mismatch often presents a formidable obstacle, impeding the efficacy of ML in scenarios characterized by small datasets, which are common in practical applications.[14–16] Bridging this divide necessitates innovative strategies and methodologies tailored to the unique challenges posed by small datasets. In this paper, we present a transfer learning strategy to address this pressing need.

In transfer learning, the knowledge learned in one setting is used to improve performance on a related target task.[17] In the first step, known as pre-training, an NN is trained on a large set of data that is closely related to the target task, yielding an NN with pre-optimized parameters. In the second step, termed fine-tuning, the same model is re-trained using data from the target task. Now, since the model is already close to the optimal set of parameters after the pre-training, much less data are needed during fine-tuning to obtain an accurate model for the target task. This strategy has proven useful in various fields, including image classification[18] and captioning,[19] gaming strategies,[20] and social network analysis.[21] In the physical sciences, a range of studies recently investigated the use of transfer learning to improve atomic force predictions for molecular dynamics simulations.[22–29] Moreover, some previous studies have used transfer learning to improve property predictions of molecules and materials.[30,31] For example, Su *et al.* recently presented a transfer learning strategy for predicting frontier orbital energies of organic materials based on DFT data obtained using two different functionals for generating pre-training and fine-tuning labels, respectively.[32] Furthermore, several studies have explored the use of transfer learning for regression on both HOPV[11,30,33–35] and FreeSolv.[30,36,37] For such tasks, Sun *et al.* found that supervised pre-training generally helps if the pre-training labels are closely aligned with the downstream tasks, but the success of efficient transfer learning strongly depends on the employed datasets, ML models, and their hyperparameters.[38] Hoffmann *et al.* further showed that increasing both the size of the pre-training and the fine-tuning datasets linearly improved the prediction results in the transfer learning settings investigated by them.[39] However, as pointed out by Dral, relatively little research is performed in exploring the advantages and limitations of transfer learning in the context of repurposing the models for different levels and molecules.[13]

In this study, we start by undertaking rigorous benchmarking, evaluating standard ML techniques in the context of small molecular datasets. In particular, we compare the use of SOAP descriptors tailored to a standard ML algorithm (KRR, NN, or gradient boosting) to the message passing neural network PaiNN for regression. The comparison is performed both on the Harvard Organic Photovoltaics (HOPV) dataset,[40] containing photovoltaic properties of 350 organic molecules, and on the FreeSolv dataset,[41] which contains solvation energies of 643 small organic molecules.

Second, we propose a transfer learning strategy as a means of augmenting predictive modeling with small datasets. It involves leveraging computations from computationally cheap *ab initio* or semi-empirical models to generate pre-training labels for PaiNN.

II. METHODS

In this study, we use two small common physical chemistry datasets, the Harvard Organic Photovoltaics (HOPV) dataset[40] and the FreeSolv dataset.[41] The HOPV dataset contains 350 molecules and their photovoltaic data obtained from both experiments and quantum-chemical calculations. In this study, we focus on the HOMO–LUMO gaps obtained from PBE0/def2-SVP density functional theory, which is a common *ab initio* method used for standard chemistry datasets.[7,42–47] Note that the version of the HOPV dataset used by us (HOPV_15_revised_2) contains a single molecule with an Se atom, which we regarded as a structural outlier and, therefore, removed from the dataset prior to training. The FreeSolv dataset contains experimental and calculated solvation energies of 643 small molecules,[41] from which we use the values obtained from "alchemical" calculations based on molecular dynamics simulations by the Mobley group.[48]

For benchmarking standard ML methods with HOPV and FreeSolv, we employed the PaiNN implementation of the *schnetpack* Python library.[49] PaiNN was recently proposed as an extension of the popular MPNN SchNet.[6] In the PaiNN architecture, molecular atoms are represented as graph nodes, and interactions between atoms are represented as undirected edges between the nodes. Each atom (node) is described by atomic scalar features s, and each interaction (edge) as edge vector features h that are updated using stacked message-passing and update blocks. After the message phasing phase, the final predictions are obtained after summing or averaging over the final node descriptors.[6]

We further used the SOAP implementation of the *dscribe* library[50] in junction with either *sklearn*'s GradientBoostingRegressor (GBoost), kernel ridge regression, or a simple feed-forward NN. The GBoost algorithm profited from an additional set of simple molecular descriptors concatenated to the SOAP descriptor. For NN learning, to limit the size of the NNs, the sparse SOAP descriptors were compressed by principal component analysis, retaining 99.999% of the variance.[7,9] The hyperparameters of PaiNN were optimized using an iterative grid search, and the hyperparameters of the SOAP-ML models were optimized using Bayesian optimization. For PaiNN, a learning rate decay and early stopping were employed, and results were averaged over five runs. In all cases, 20% of the data was used for testing, and for the NNs, 20% of the data was used for validation and early stopping. The final model hyperparameters are listed in the Appendix.

In the second part of this study, we present a method to improve the predictions of the message passing neural network PaiNN,[6] which was also identified as the best ML model for the datasets at hand (see Sec. III), via transfer learning. For this, we need (i) a (small) molecular dataset with high-accuracy data of the target property (y_{true}), (ii) a large dataset of molecular structures to be used for pre-training, and (iii) a computationally inexpensive method from the field of quantum chemistry to compute the target property for the pre-training data. Cheap *ab initio* methods, like density functional theory in the linear density approximation (LDA-DFT),

06 January 2025 06:56:56

© Author(s) 2024

or semi-empirical methods typically yield results that are qualitatively correct but quantitatively far off.[51] To correct for the biases inherent in the results obtained from these methods, we normalize both datasets to have mean 0 and standard deviation 1 before NN training.

In this study, we use the molecular structures of the OE62 dataset as pre-training structures for the HOPV dataset, and the molecular structures of the QM9 dataset as pre-training structures for FreeSolv (ii). The OE62 dataset contains crystal structures of 62 k medium-sized organic molecules,[47] and the QM9 dataset contains DFT-optimized structures of 134 k small molecules.[52] To obtain pre-training labels for these datasets (iii), HOMO–LUMO gaps and solvation energies were calculated with LDA-DFT using the ORCA software[53] and tight-binding density functional theory (DFTB) using the GFN2-xTB method[54] implemented in the XTB software.[55] Solvation energies were obtained as the difference in single point energy of the molecule in vacuum and the molecule with an implicit solvent model for water (conductor-like polarizable continuum model, CPCM in ORCA,[56] and generalized Born model with surface area contributions, GBSA in XTB,[57–59] respectively). Note that we do not use the properties originally published with the respective datasets QM9 and OE62 during pre-training. To illustrate the data quality obtained from LDA and XTB computations, we display the scatter plots of the original HOPV and FreeSolv data, respectively, against the data obtained from LDA and XTB, respectively, corrected by linear regression (Fig. 1). These plots clearly show that the obtained HOMO–LUMO gaps are very close to the original HOPV data, but the solvation energies are quite dissimilar to the original FreeSolv data. Furthermore, for both datasets, the properties obtained from LDA are closer to the original data than those obtained from XTB. The mean absolute errors of the fits are 0.05 eV (HOPV LDA), 0.12 eV (HOPV XTB), 1.1 kcal/mol (FreeSolv LDA), and 1.7 kcal/mol (FreeSolv XTB).

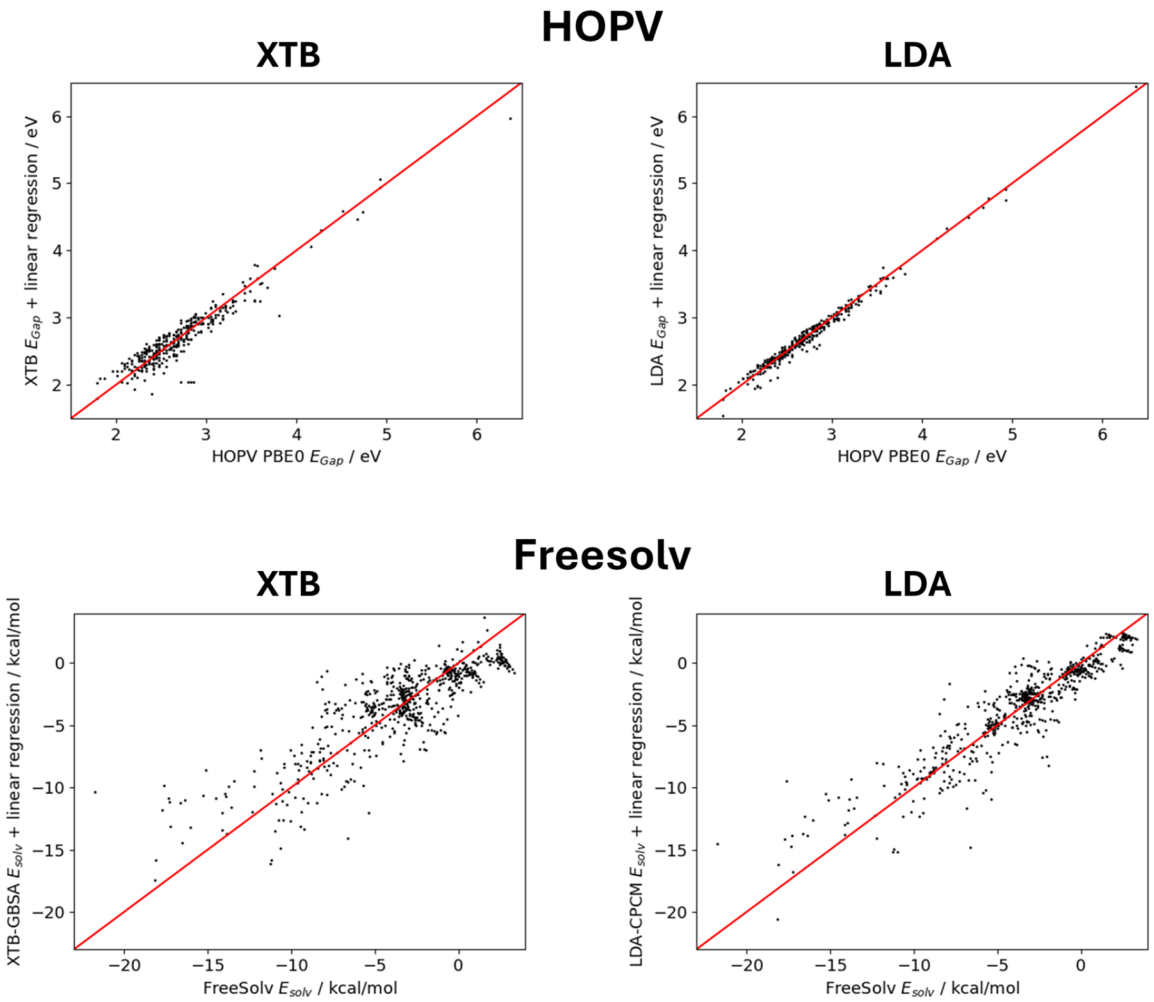


FIG. 1. Scatter plots of the original data plotted against data obtained from XTB (left) or LDA (right) plus linear regression for the datasets HOPV (HOMO–LUMO-gaps, top) and FreeSolv (solvation energies, bottom).

06 January 2025 06:56:56

© Author(s) 2024

On average, HOMO–LUMO gap computations took ~106 s (LDA-DFT) and 0.22 s (XTB) per data point in OE62, respectively. Solvation energy computations took ~70 s (LDA-DFT) and 0.15 s (XTB) per data point in QM9, respectively, and 1.33 s (XTB) per data point in OE62.

Pre-training on QM9 data was performed for 200 epochs (100 k training, 10 k validation, 20 k test data points). Pre-training on OE62 data was performed for 250 epochs for LDA and XTB data (50 k training, 5 k validation, 6 k test data points). For each method, three models were trained with different random seeds, and the best model as evaluated on the test subset was selected for fine-tuning.

III. RESULTS AND DISCUSSION

A. ML benchmarks

We start by benchmarking standard chemistry ML models for regression on the two small, but high-quality datasets under investigation, HOPV[40] and FreeSolv,[41] listed under (i) in the Methods section. The results are summarized in Table I as mean absolute errors [MAEs, root mean squared errors (RMSE) are given in parenthesis]. For both datasets, very low errors are obtained with PaiNN, yielding MAEs of 0.01 eV for HOPV and 0.56 kcal/mol for FreeSolv, respectively. For HOPV, the SOAP + SD + GBoost model yields an MAE of 0.19 eV, while 0.86 kcal/mol is obtained for regression on FreeSolv. The results obtained with these algorithms are qualitatively similar to those obtained in previous ML studies on HOPV[11,30,33–35,60,61] and FreeSolv.[30,36,37,62–70] The remaining models SOAP + NN and SOAP + KRR perform considerably worse on both datasets, in direct comparison, SOAP + NN yields better results for HOPV, while SOAP + KRR yields better results for FreeSolv.

B. Transfer learning

We use cheap models from quantum chemistry (LDA-DFT and XTB-DFTB) normalized to the same mean (zero) and standard deviation (one) as the fine-tuning data to obtain labels for pre-training the PaiNN models (details in the Methods section). For HOPV, consisting of medium-sized organic molecules, we use the structures of the OE62 dataset for pre-training, which contains ~62 k diverse structures of medium-sized organic molecules. For FreeSolv, which on the other hand consists of small organic molecules, we use the structures of the QM9 dataset for pre-training, which contains ~132 k organic molecules with up to nine heavy (non-H) atoms. Tables II and III summarize the results obtained for regression on the HOPV and FreeSolv dataset, respectively, after PaiNN training from scratch (cf. Table I) or after pre-training on data labeled with either XTB or LDA, respectively. The learning curves displaying the MAE vs the number of fine-tuning training points are shown in Fig. 2 (here, pre-training is performed on the full respective pre-training dataset), and the learning curves displaying the MAE vs the number of pre-training points are shown in Fig. 3 (here, fine-tuning is performed on the full respective fine-tuning dataset), each for the HOPV set on top and the FreeSolv set in the bottom. All reported MAEs are obtained as the mean over five independent runs, and the error bars indicate the MAEs' standard deviations.

For HOPV (Table II), pre-training results in an improved learning performance. The MAE is reduced from 7.6 meV (training from scratch) to 6.3 meV after pre-training on OE62 + XTB and further to 6.1 meV eV after pre-training on OE62 + LDA. The fine-tuning learning curves for HOPV (Fig. 2, top) show that this improvement of ~1.5 meV over training from scratch (black curve) after pre-training on OE62 + XTB (red) or OE62 + LDA (blue) is relatively stable, independent of the number of HOPV fine-tuning training points. In addition, the standard deviation of the MAEs is reduced after pre-training on either dataset.

The HOPV pre-training learning curves (Fig. 3, top) confirm that for all pre-training dataset sizes tested, the labels obtained from the QC models lead to a higher accuracy as compared to training from scratch. Interestingly, increasing the size of the pre-training dataset does not monotonically improve the final training results. Instead, the final training errors after pre-training on 10k data points do not strongly improve when more pre-training data (up to 50k

TABLE I. Mean absolute errors for HOPV HOMO–LUMO gaps (in eV) and FreeSolv solvation energies (in kcal/mol) from four different models. The first three models use SOAP molecular descriptors connected to either KRR, a dense NN, or GBoost. The GBoost model uses an additional set of simple molecular descriptors (SD) concatenated with the SOAP descriptors. The final model is the message passing NN PaiNN. Uncertainties are measured in RMSE (in parentheses). Best MAE results are bold.

Model	HOPV MAE (RMSE) (eV)	FreeSolv MAE (RMSE) (kcal/mol)
SOAP + KRR	0.31 (0.46)	1.7 (2.2)
SOAP + NN	0.25 (0.33)	2.1 (3.2)
SOAP + SD + GBoost	0.19 (0.30)	0.86 (1.30)
PaiNN	**0.01 (0.01)**	**0.56 (0.86)**

TABLE II. MAEs for HOPV HOMO–LUMO gaps (in eV) from PaiNN obtained after training from scratch or after pre-training on OE62 data with labels obtained from XTB and LDA, respectively. Uncertainties are measured in RMSE (in parentheses). Best MAE results are bold.

Pre-training	HOPV MAE (RMSE) (meV)
None	7.6 (11.9)
OE62+XTB	6.3 (9.7)
OE62+LDA	**6.1 (9.6)**

TABLE III. MAEs for FreeSolv solvation energies (in kcal/mol) from PaiNN obtained after training from scratch or after pre-training on OE62 data with labels obtained from XTB and LDA, respectively. Uncertainties are measured in RMSE (in parentheses). Best MAE results are bold.

Pre-training	FreeSolv MAE (RMSE) (kcal/mol)
None	**0.56 (0.86)**
QM9+XTB	0.61 (1.05)
QM9+LDA	0.66 (1.10)

06 January 2025 06:56:56

© Author(s) 2024

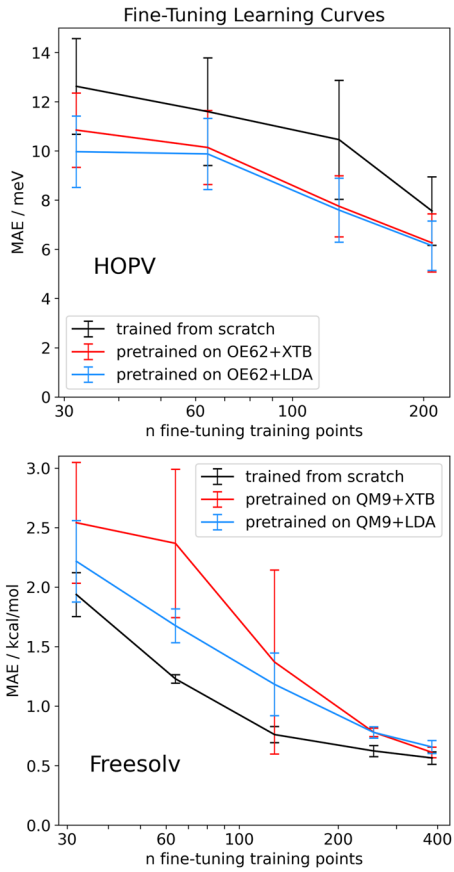


FIG. 2. Fine-tuning learning curves for PaiNN training on HOPV (top) and FreeSolv (bottom), after training from scratch (black) or after pre-training on OE62 or QM9 data, respectively, with labels obtained from XTB (red) and LDA-DFT (blue). The MAE (mean and standard deviation over five runs) is plotted against the number of training examples used for fine-tuning (log scale).

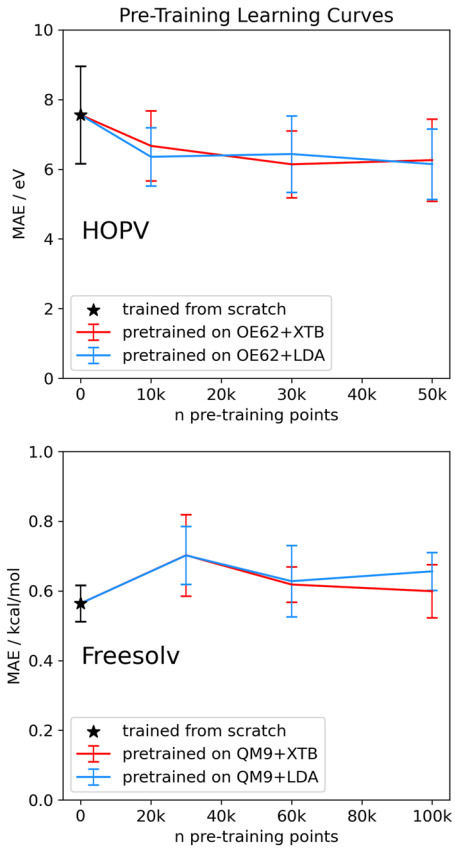


FIG. 3. Pre-training learning curves for PaiNN training on HOPV (top) and FreeSolv (bottom), after training from scratch (black star) or after pre-training on OE62 or QM9 data, respectively, with labels obtained XTB (red) and LDA-DFT (blue). The MAE (mean and standard deviation over five runs) is plotted against the number of training examples used for pre-training.

data points) is used. This indicates that a pre-trained model that is too strongly fitted to the pre-training task may be worse for transfer learning than a more biased model that only learned to grasp high-level concepts from a limited amount of training data, which contradicts the findings in the recent study of Hoffmann *et al.*[39]

Regression learning on FreeSolv, on the other hand, is not improved by pre-training the PaiNN models (Table III). Instead, the final training MAEs are increased from 0.56 kcal/mol (training from scratch) to 0.61 kcal/mol (XTB) and 0.66 kcal/mol (LDA). The fine-tuning learning curves (Fig. 2, bottom) show that no improvement is achieved throughout, independent of the size of the pre-training dataset. In addition, there is no clear trend in how the size of the pre-training dataset influences the result of the subsequent fine-tuning. In addition, the MAE standard deviations increase when pre-training on either dataset.

As discussed before, the regression task on FreeSolv might be more complex as compared to regression on HOPV. The methods used for producing the respective pre-training labels (QC methods) are quite similar to the method used for generating HOPV data (standard DFT),[40] but quite dissimilar to the method used for generating FreeSolv data ("alchemical" calculations based on molecular dynamics simulations; see also the scatter plots in Fig. 1).[48] These factors may be the reason for why the employed transfer learning strategy does work very well for use with the HOPV dataset, but not for FreeSolv.

Another hypothesis for why transfer learning might fail in certain scenarios is that the underlying regression task is too easy and, therefore, no additional pre-training is needed.[38] However, previous results for regression on molecular energies on the QM9 dataset (free energy G, internal energy U, enthalpy H) using PaiNN indicated that MAEs as low as ~0.14 kcal/mol can be obtained on a sufficiently large dataset, which is considerably lower than the MAEs obtained here.[6] In addition, these quantities (G, U, H) can easily be calculated from only the respective molecular structure, but solvation energy computations additionally need to take into account the shape and intermolecular interactions of the respective solvation shell, thus adding another level of complexity. Therefore, we anticipate that the regression task on FreeSolv is not too easy for transfer learning.

06 January 2025 06:56:56

© Author(s) 2024

In an attempt to obtain pre-trained models that are more robust and, therefore, allegedly more suitable for transfer learning, we added additional data points to the respective pre-training datasets. For this, we used XTB to generate additional labels for QM9 (HOMO–LUMO gaps) and OE62 (solvation energies). However, we found that pre-training on a joint OE62 and QM9 dataset did not improve the final fine-tuning results as compared to pre-training on OE62 or QM9 data only, respectively, for both HOPV and FreeSolv training. This is surprising because NNs usually get more accurate when increasing the amount of training data, and new dissimilar data points can help to create more robust models. Here, instead, the introduction of more data points that are dissimilar to the data contained in the final fine-tuning dataset apparently impedes the learning process. In particular, for FreeSolv, the underlying problem of having a complex learning task and dissimilar methods for generating pre-training and fine-tuning labels could not be solved by augmenting the pre-training dataset. This observation is in line with previous work on transfer learning in chemistry, which stated that the pre-training and fine-tuning datasets should be as closely aligned as possible to obtain the best results.[30]

Following this idea, we aimed to instead make the pre-training and fine-tuning datasets more similar to each other: The HOPV dataset consists almost entirely of S-functionalized aromatic molecules, while the OE62 dataset used for pre-training contains a much more diverse set of molecules, including more structural variety and a wider range of atomic species. Accordingly, we filtered all molecules from the OE62 pre-training dataset containing any of the elements As, Se, Br, Te, or I, and furthermore removed all molecules that do not contain any sulfur atoms, thereby reducing the dataset size from 61 489 to 41 487 molecules. Pre-training with 20 k less molecules from the filtered OE62 dataset resulted in a slight loss of performance compared to pre-training of the full OE62 + XTB dataset (+0.01 eV MAE), which is approximately the same result as obtained after non-selective filtering of the pre-training dataset (+0.02 eV MAE, cf. Fig. 3). This indicates that increased similarity between pre-training and fine-tuning structures might slightly improve the learning performance by yielding more tailored pre-trained models, but the effect found here is very small.

Finally, previous studies have used discriminative fine-tuning to further optimize transfer learning.[29] In this method, the learning rate during fine-tuning is adapted layer-wise, such that the first layers have a much smaller learning rate than the final layers of the model. The idea is that the first layers, which learn to grasp fundamental features of the molecular data, will be sufficiently trained during pre-training, while the last layers, which derive the final predictions, will need more fine-tuning. We implemented discriminative fine-tuning for PaiNN, following the approach outlined by Zaverkin *et al.*[29] Here, the learning rate (lr) for layer n−1 was adjusted as $\mathrm{lr}_{n-1} = \mathrm{lr}_n/5$, with $\mathrm{lr} = 5 \times 10^{-4}$ being the learning rate of the last layer. However, we found that this method did not yield any improvements over non-discriminative fine-tuning for transfer learning with PaiNN. In contrast, we found that in general, modifications of the learning rate and learning rate decay had a significant impact on the learning performance, and sticking to the parameters identified by the initial hyperparameter sweep continued to yield the best results.

IV. CONCLUSIONS

This study investigated the use of common ML algorithms for use with small datasets, focusing on HOPV (350 data points) and FreeSolv (643 data points). We found that for these datasets, both the MPNN PaiNN and gradient boosting with regression trees operating on SOAP molecular descriptors concatenated to a set of simple molecular descriptors yield accurate results, clearly outperforming the SOAP + NN and SOAP + KRR architectures.

Furthermore, we presented a method for data-efficient ML on small physical chemistry datasets using transfer learning, which we tested on the two datasets HOPV and FreeSolv. The pre-training data were obtained from computationally cheap QC calculations corrected by simple linear regression calibrated on the fine-tuning dataset. This method utilizes the fact that many cheap *ab initio* or semi-empirical techniques yield quantitatively incorrect but qualitatively correct data, which can easily be corrected via linear regression. For the HOPV dataset, the introduced method outperforms transfer learning with labels obtained from an ML model. Here, pre-training with labels computed by either XTB or LDA results in very similar learning performances, but XTB computations were much faster than LDA computations. Accordingly, labeling the pre-training data with XTB should be preferred.

However, the same method did not yield improved learning results for regression on the FreeSolv dataset. We suppose that this may be due to the fact that the underlying learning task is much more complex and the data used for pre-training are too dissimilar to the data contained in FreeSolv.

Furthermore, the ML results obtained for HOPV indicate that pre-training with a limited amount of data points may be beneficial since the obtained models are slightly more biased and less strongly fitted to the pre-training task. Furthermore, we found that augmenting the pre-training datasets with additional structures that are less similar to the fine-tuning data does not improve the learning performance. Finally, for the models trained here, sticking to the hyperparameters identified by the initial hyperparameter sweeps was more efficient than adapting the learning rates as performed in discriminative fine-tuning. However, given the limited scope of this study, further research will be needed to investigate whether these findings can be generalized to other datasets as well.

In conclusion, the results obtained from this exploratory study suggest that for chemistry ML tasks that use small datasets, an effective transfer learning strategy should use pre-training on molecules that are similar to the molecules of the downstream task, a labeling method that is closely aligned to the method used for labeling the fine-tuning data, and an optimized training set size for the pre-training. Future studies may investigate whether these findings apply to other datasets as well.

ACKNOWLEDGMENTS

T.K. and A.B. acknowledge support from the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS). Computing resources were kindly provided by the Freie Universität Berlin hpc cluster Curta[71] and by the Helmholtz-Zentrum Dresden-Rossendorf. The introduction was in part formulated with the help of ChatGPT 3.5.

06 January 2025 06:56:56

© Author(s) 2024

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Thorren Kirschbaum: Conceptualization (lead); Data curation (lead); Formal analysis (lead); Funding acquisition (supporting); Investigation (lead); Methodology (lead); Project administration (supporting); Software (lead); Validation (equal); Visualization (lead); Writing – original draft (lead); Writing – review & editing (supporting). **Annika Bande**: Investigation (supporting); Project administration (lead); Resources (lead); Supervision (lead); Validation (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are openly available at https://github.com/ThorrenKirschbaum/TransferLearningPaiNN.

APPENDIX: MACHINE LEARNING MODEL HYPERPARAMETERS

We here provide details of the ML models used for regression on both HOPV and FreeSolv (cf. Table I).

The final SOAP + KRR models used the following hyperparameters:

HOPV: $r_{cut}^{SOAP} = 5.63$, $n_{\max}^{SOAP} = 7$, $l_{\max}^{SOAP} = 7$, $alpha_{KRR} = 1.12$, kernel: cosine.

FreeSolv: $r_{cut}^{SOAP} = 9.53$, $n_{\max}^{SOAP} = 8$, $l_{\max}^{SOAP} = 8$, $alpha_{KRR} = 0.59$, kernel: = cosine.

In the SOAP + NN model, the neural network consists of densely connected layers with steadily decreasing numbers of neurons, including batch normalization and a dropout layer with $p = 0.3$ after the first layer. Here, 20% of the training data was used for validation, and early stopping was employed. The final models used the following hyperparameters:

HOPV: $r_{cut}^{SOAP} = 7.22$, $n_{\max}^{SOAP} = 8$, $l_{\max}^{SOAP} = 8$, $n_{layers} = 2$, $lr = 0.008\,60$, nonlinearity: swish.

FreeSolv: $r_{cut}^{SOAP} = 9.94$, $n_{\max}^{SOAP} = 7$, $l_{\max}^{SOAP} = 7$, $n_{layers} = 4$, $lr = 0.006\,27$, nonlin.: swish.

In the SOAP + SD + GBoost model, the following simple molecular descriptors were employed: total number of atoms, number of atoms for each atom type in the dataset, one-hot-encoding for each atom type (present/not present in the respective molecule), mean and standard deviation of the carbon–carbon-distance of neighboring carbon atoms within the molecule. Using these descriptors with the KRR or NN did not yield any improvement for the respective learning algorithm. The final SOAP + SD + GBoost models used the following hyperparameters:

HOPV: $r_{cut}^{SOAP} = 9.71$, $n_{\max}^{SOAP} = 5$, $l_{\max}^{SOAP} = 8$, $lr_{GBoost} = 0.059$, $n_estimators_{GBoost} = 295$, $\max_depth_{GBoost} = 4$.

FreeSolv: $r_{cut}^{SOAP} = 6.85$, $n_{\max}^{SOAP} = 4$, $l_{\max}^{SOAP} = 4$, $lr_{GBoost} = 0.250$, $n_estimators_{GBoost} = 158$, $\max_depth_{GBoost} = 4$.

The final PaiNN models used the following hyperparameters:

HOPV: *radial_basis* = *GaussianRBF*, $r_{cut} = 5.0$, *n_rbf* = 20, *n_atom_basis* = 30, *n_interactions* = 3, *cutoff_fn* = *CosineCutoff*, $lr = 5e-4$, *lr_decay_factor* = 0.5, *lr_decay_patience* = 5, *batchsize* = 10.

FreeSolv: *radial_basis* = *GaussianRBF*, $r_{cut} = 8.0$, *n_rbf* = 30, *n_atom_basis* = 200, *n_interactions* = 3, *cutoff_fn* = *CosineCutoff*, $lr = 5e-4$, *lr_decay_factor* = 0.5, *lr_decay_patience* = 5, *batchsize* = 10.

REFERENCES

[1] G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, and O. Anatole von Lilienfeld, "Machine learning of molecular electronic properties in chemical compound space," New J. Phys. **15**, 095003 (2013).

[2] K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, "Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space," J. Phys. Chem. Lett. **6**, 2326–2331 (2015).

[3] P. O. Dral, "Quantum chemistry in the age of machine learning," J. Phys. Chem. Lett. **11**, 2336–2347 (2020).

[4] V. Fung, J. Zhang, E. Juarez, and B. G. Sumpter, "Benchmarking graph neural networks for materials chemistry," npj Comput. Mater. **7**, 84 (2021).

[5] K. Singh, J. Münchmeyer, L. Weber, U. Leser, and A. Bande, "Graph neural networks for learning molecular excitation spectra," J. Chem. Theory Comput. **18**, 4408–4417 (2022).

[6] K. Schütt, O. Unke, and M. Gastegger, "Equivariant message passing for the prediction of tensorial properties and molecular spectra," in *Proceedings of the 38th International Conference on Machine Learning* (PMLR, 2021), pp. 9377–9388.

[7] T. Kirschbaum, B. von Seggern, J. Dzubiella, A. Bande, and F. Noé, "Machine learning frontier orbital energies of nanodiamonds," J. Chem. Theory Comput. **19**, 4461–4473 (2023).

[8] A. P. Bartók, R. Kondor, and G. Csányi, "On representing chemical environments," Phys. Rev. B **87**, 184115 (2013).

[9] B. Parsaeifard, D. S. De, A. S. Christensen, F. A. Faber, E. Kocer, S. De, J. Behler, O. A. v. Lilienfeld, and S. Goedecker, "An assessment of the structural resolution of various fingerprints commonly used in machine learning," Mach. Learn.: Sci. Technol. **2**, 015018 (2021).

[10] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, and K.-R. Müller, "Machine learning force fields," Chem. Rev. **121**, 10142–10186 (2021).

[11] Z.-W. Zhao, M. del Cueto, Y. Geng, and A. Troisi, "Effect of increasing the descriptor set on machine learning prediction of small molecule-based organic solar cells," Chem. Mater. **32**, 7777–7787 (2020).

[12] Z.-W. Zhao, Y. Geng, A. Troisi, and H. Ma, "Performance prediction and experimental optimization assisted by machine learning for organic photovoltaics," Adv. Intell. Syst. **4**, 2100261 (2022).

[13] P. O. Dral, "AI in computational chemistry through the lens of a decade-long journey," Chem. Commun. **60**, 3240–3258 (2024).

[14] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, "Big data meets quantum chemistry approximations: The Δ-machine learning approach," J. Chem. Theory Comput. **11**, 2087–2096 (2015).

[15] Y. Zhang and C. Ling, "A strategy to apply machine learning to small datasets in materials science," npj Comput. Mater. **4**, 25 (2018).

[16] B. Mazouin, A. A. Schöpfer, and O. A. von Lilienfeld, "Selected machine learning of HOMO–LUMO gaps with improved data-efficiency," Mater. Adv. **3**, 8306–8316 (2022).

[17] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," Proc. IEEE **109**, 43–76 (2020).

06 January 2025 06:56:56

© Author(s) 2024

[18]K. Saenko, B. Kulis, M. Fritz, and T. Darrell, K. Daniilidis, P. Maragos, and N. Paragios (Eds.) in *Computer Vision-ECCV 2010* (Springer Berlin, Heidelberg, 2010).
[19]O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2015), pp. 3156–3164.
[20]M. Sharma, M. Holmes, J. Santamaria, A. Irani, and J. Isbell, "Transfer learning in real-time strategy games using hybrid CBR/RL," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence* (Morgan Kaufmann Publishers, Inc., 2007), pp. 1041–1046.
[21]J. Tang, T. Lou, J. Kleinberg, and S. Wu, "Transfer learning to infer social ties across heterogeneous networks," ACM Trans. Inf. Syst. **34**, 1–43 (2016).
[22]J. S. Smith, B. T. Nebgen, R. Zubatyuk, N. Lubbers, C. Devereux, K. Barros, S. Tretiak, O. Isayev, and A. E. Roitberg, "Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning," Nat. Commun. **10**, 2903 (2019).
[23]S. Käser, D. Koner, A. S. Christensen, O. A. von Lilienfeld, and M. Meuwly, J. Phys. Chem. A **124**, 8853–8865 (2020).
[24]P. Zheng, R. Zubatyuk, W. Wu, O. Isayev, and P. O. Dral, "Artificial intelligence-enhanced quantum chemical method with broad applicability," Nat. Commun. **12**, 7022 (2021).
[25]S. Käser, J. O. Richardson, and M. Meuwly, "Transfer learning for affordable and high-quality tunneling splittings from instanton calculations," J. Chem. Theory Comput. **18**, 6840–6850 (2022).
[26]X. Gao, W. Gao, W. Xiao, Z. Wang, C. Wang, and L. Xiang, "Supervised pre-training for molecular force fields and properties prediction," arXiv:2211.14429 (2022).
[27]D. Zhang, H. Bi, F.-Z. Dai, W. Jiang, X. Liu, L. Zhang, and H. Wang, "Pretraining of attention-based deep learning potential model for molecular simulation," npj Comput. Mater. **10**, 94 (2024).
[28]M. S. Chen, J. Lee, H.-Z. Ye, T. C. Berkelbach, D. R. Reichman, and T. E. Markland, "Data-efficient machine learning potentials from transfer learning of periodic correlated electronic structure methods: Liquid water at AFQMC, CCSD, and CCSD(T) accuracy," J. Chem. Theory Comput. **19**, 4510–4519 (2023).
[29]V. Zaverkin, D. Holzmüller, L. Bonfirraro, and J. Kästner, "Transfer learning for chemically accurate interatomic neural network potentials," Phys. Chem. Chem. Phys. **25**, 5383–5396 (2023).
[30]H. Li, X. Zhao, S. Li, F. Wan, D. Zhao, and J. Zeng, "Improving molecular property prediction through a task similarity enhanced transfer learning strategy," iScience **25**, 105231 (2022).
[31]Z. Shui, D. Karls, M. Wen, E. Tadmor, G. Karypis *et al.*, "Injecting domain knowledge from empirical interatomic potentials to neural networks for predicting material properties," Adv. Neural Inf. Process. Syst. **35**, 14839–14851 (2022).
[32]A. Su, X. Zhang, C. Zhang, D. Ding, Y.-F. Yang, K. Wang, and Y.-B. She, "Deep transfer learning for predicting frontier orbital energies of organic materials using small data and its application to porphyrin photocatalysts," Phys. Chem. Chem. Phys. **25**, 10536 (2023).
[33]A. Eibeck, D. Nurkowski, A. Menon, J. Bai, J. Wu, L. Zhou, S. Mosbach, J. Akroyd, and M. Kraft, "Predicting power conversion efficiency of organic photovoltaics: Models and data analysis," ACS Omega **6**, 23764–23775 (2021).
[34]J. Munshi, W. Chen, T. Chien, and G. Balasubramanian, "Transfer learned designer polymers for organic solar cells," J. Chem. Inf. Model. **61**, 134–142 (2021).
[35]G. J. Moore, O. Bardagot, and N. Banerji, "Deep transfer learning: A fast and accurate tool to predict the energy levels of donor molecules for organic photovoltaics," Adv. Theory Simul. **5**, 2100511 (2022).
[36]F. H. Vermeire and W. H. Green, "Transfer learning for solvation free energies: From quantum chemistry to experiments," Chem. Eng. J. **418**, 129307 (2021).
[37]D. Zhang, S. Xia, and Y. Zhang, "Accurate prediction of aqueous free solvation energies using 3D atomic feature-based graph neural network with transfer learning," J. Chem. Inf. Model. **62**, 1840–1848 (2022).
[38]R. Sun, H. Dai, and A. W. Yu, "Does GNN pretraining help molecular representation?," Adv. Neural Inf. Process. Syst. **35**, 12096–12109 (2022).
[39]N. Hoffmann, J. Schmidt, S. Botti, and M. A. L. Marques, "Transfer learning on large datasets for the accurate prediction of material properties," Digital Discovery **2**, 1368–1379 (2023).
[40]S. A. Lopez, E. O. Pyzer-Knapp, G. N. Simm, T. Lutzow, K. Li, L. R. Seress, J. Hachmann, and A. Aspuru-Guzik, "The Harvard organic photovoltaic dataset," Sci. Data **3**, 160086 (2016).
[41]G. Duarte Ramos Matos, D. Y. Kyu, H. H. Loeffler, J. D. Chodera, M. R. Shirts, and D. L. Mobley, "Approaches for calculating solvation free energies and enthalpies demonstrated with an update of the FreeSolv database," J. Chem. Eng. Data **62**, 1559–1569 (2017).
[42]C. Adamo and V. Barone, "Toward reliable density functional methods without adjustable parameters: The PBE0 model," J. Chem. Phys. **110**, 6158–6170 (1999).
[43]J. P. Perdew, M. Ernzerhof, and K. Burke, "Rationale for mixing exact exchange with density functional approximations," J. Chem. Phys. **105**, 9982–9985 (1996).
[44]A. Schäfer, H. Horn, and R. Ahlrichs, "Fully optimized contracted Gaussian basis sets for atoms Li to Kr," J. Chem. Phys. **97**, 2571–2577 (1992).
[45]F. Weigend and R. Ahlrichs, "Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy," Phys. Chem. Chem. Phys. **7**, 3297–3305 (2005).
[46]M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," Phys. Rev. Lett. **108**, 058301 (2012).
[47]A. Stuke, C. Kunkel, D. Golze, M. Todorović, J. T. Margraf, K. Reuter, P. Rinke, and H. Oberhofer, "Atomic structures and orbital energies of 61,489 crystal-forming organic molecules," Sci. Data **7**, 58 (2020).
[48]D. L. Mobley and J. P. Guthrie, "FreeSolv: A database of experimental and calculated hydration free energies, with input files," J. Comput. Aided Mol. Des. **28**, 711–720 (2014).
[49]K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko, and K.-R. Müller, "SchNetPack: A deep learning toolbox for atomistic systems," J. Chem. Theory Comput. **15**, 448–455 (2019).
[50]L. Himanen, M. O. J. Jäger, E. V. Morooka, F. Federici Canova, Y. S. Ranawat, D. Z. Gao, P. Rinke, and A. S. Foster, "DScribe: Library of descriptors for machine learning in materials science," Comput. Phys. Commun. **247**, 106949 (2020).
[51]W. Thiel, "Semiempirical quantum–chemical methods," Wiley Interdiscip. Rev.: Comput. Mol. Sci. **4**, 145–157 (2014).
[52]R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, "Quantum chemistry structures and properties of 134 kilo molecules," Sci. Data **1**, 140022 (2014).
[53]F. Neese, F. Wennmohs, U. Becker, and C. Riplinger, "The ORCA quantum chemistry program package," J. Chem. Phys. **152**, 224108 (2020).
[54]C. Bannwarth, S. Ehlert, and S. Grimme, "GFN2-xTB—An accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions," J. Chem. Theory Comput. **15**, 1652–1671 (2019).
[55]C. Bannwarth, E. Caldeweyher, S. Ehlert, A. Hansen, P. Pracht, J. Seibert, S. Spicher, and S. Grimme, "Extended tight-binding quantum chemistry methods," Wiley Interdiscip. Rev.: Comput. Mol. Sci. **11**, e1493 (2021).
[56]V. Barone and M. Cossi, "Quantum calculation of molecular energies and energy gradients in solution by a conductor solvent model," J. Phys. Chem. A **102**, 1995–2001 (1998).
[57]A. Onufriev, D. Bashford, and D. A. Case, "Exploring protein native states and large-scale conformational changes with a modified generalized born model," Proteins **55**, 383–394 (2004).
[58]G. Sigalov, A. Fenley, and A. Onufriev, "Analytical electrostatics for biomolecules: Beyond the generalized Born approximation," J. Chem. Phys. **124**, 124902 (2006).
[59]A. W. Lange and J. M. Herbert, "Improving generalized born models by exploiting connections to polarizable continuum models. I. An improved effective coulomb operator," J. Chem. Theory Comput. **8**, 1999–2011 (2012).
[60]A. Paul, A. Furmanchuk, W.-k. Liao, A. Choudhary, and A. Agrawal, "Property prediction of organic donor molecules for photovoltaic applications using extremely randomized trees," Mol. Inf. **38**, 1900038 (2019).

06 January 2025 06:56:56

© Author(s) 2024

[61] C. Fare, P. Fenner, M. Benatan, A. Varsi, and E. O. Pyzer-Knapp, "A multi-fidelity machine learning approach to high throughput materials screening," npj Comput. Mater. **8**, 257 (2022).
[62] S. T. Hutchinson and R. Kobayashi, "Solvent-specific featurization for predicting free energies of solvation through machine learning," J. Chem. Inf. Model. **59**, 1338–1346 (2019).
[63] J. Scheen, W. Wu, A. S. J. S. Mey, P. Tosco, M. Mackey, and J. Michel, "Hybrid alchemical free energy/machine-learning methodology for the computation of hydration free energies," J. Chem. Inf. Model. **60**, 5331–5339 (2020).
[64] J. Weinreich, N. J. Browning, and O. A. von Lilienfeld, "Machine learning of free energies in chemical compound space using ensemble representations: Reaching experimental uncertainty for solvation," J. Chem. Phys. **154**, 134113 (2021).
[65] D. Borgis, S. Luukkonen, L. Belloni, and G. Jeanmairet, "Accurate prediction of hydration free energies and solvation structures using molecular density functional theory with a simple bridge functional," J. Chem. Phys. **155**, 024117 (2021).
[66] P. Gao, X. Yang, Y.-H. Tang, M. Zheng, A. Andersen, V. Murugesan, A. Hollas, and W. Wang, "Graphical Gaussian process regression model for aqueous solvation free energy prediction of organic molecules in redox flow batteries," Phys. Chem. Chem. Phys. **23**, 24892–24904 (2021).
[67] Y. Pathak, S. Mehta, and U. D. Priyakumar, "Learning atomic interactions through solvation free energy prediction using graph neural networks," J. Chem. Inf. Model. **61**, 689–698 (2021).
[68] Y. Chung, F. H. Vermeire, H. Wu, P. J. Walker, M. H. Abraham, and W. H. Green, "Group contribution and machine learning approaches to predict Abraham solute parameters, solvation free energy, and solvation enthalpy," J. Chem. Inf. Model. **62**, 433–446 (2022).
[69] K. Low, M. L. Coote, and E. I. Izgorodina, "Explainable solvation free energy prediction combining graph neural networks with chemical intuition," J. Chem. Inf. Model. **62**, 5457–5470 (2022).
[70] Z.-Y. Zhang, D. Peng, L. Liu, L. Shen, and W.-H. Fang, "Machine learning prediction of hydration free energy with physically inspired descriptors," J. Phys. Chem. Lett. **14**, 1877–1884 (2023).
[71] L. Bennett, B. Melchers, and B. Proppe, *Curta: A General-Purpose High-Performance Computer at ZEDAT* (Freie Universität Berlin, Germany, 2020).

06 January 2025 06:56:56

© Author(s) 2024