

Machine learning from quantum chemistry to predict experimental solvent effects on reaction rates

Yunsie Chung and William H. Green*

Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, 02139, U.S.A

E-mail: whgreen@mit.edu

Abstract

Fast and accurate prediction of solvent effects on reaction rates are crucial for kinetic modeling, chemical process design, and high-throughput solvent screening. Despite the recent advance in machine learning, a scarcity of reliable data has hindered the development of predictive models that are generalizable for diverse reactions and solvents. In this work, we generate a large set of data with the COSMO-RS method for over 28,000 neutral reactions and 295 solvents and train a machine learning model to predict the solvation free energy and solvation enthalpy of activation ($\Delta\Delta G_{\text{solv}}^\ddagger$, $\Delta\Delta H_{\text{solv}}^\ddagger$) for a solution phase reaction. On unseen reactions, the model achieves mean absolute errors of 0.71 and 1.03 kcal/mol for $\Delta\Delta G_{\text{solv}}^\ddagger$ and $\Delta\Delta H_{\text{solv}}^\ddagger$, respectively, relative to the COSMO-RS calculations. The model also provides reliable predictions of relative rate constants within a factor of 4 when tested on experimental data. The presented model can provide nearly instantaneous predictions of kinetic solvent effects or relative rate constants for a broad range of neutral closed-shell or free radical reactions and solvents only based on atom-mapped reaction SMILES and solvent SMILES strings.

1 Introduction

Accurate prediction of reaction rates is essential for modeling a variety of chemical kinetic systems such as pyrolysis,^{1,2} polymerization,³ oxidative degradation,^{4,5} and atmospheric chemistry.⁶ Detailed kinetic models enable one to predict key products, identify major kinetic pathways, and optimize reaction conditions for complex chemical systems. Kinetic mechanisms often involve hundreds to tens of thousands of elementary reactions,⁷ and a fast, high-throughput method to estimate reaction rates is thus needed. Ab initio methods like quantum mechanics/molecular mechanics (QM/MM) can provide accurate predictions of rate constants, but their high computational cost has been a major limiting factor for large-scale, automated predictions. As more kinetic data become available, data-driven approaches such as linear group contribution,^{8–10} decision tree based rate rules,^{11,12} and machine learning (ML) models^{13–19} have emerged as more popular choices for estimating kinetic parameters. Several ML models^{15–17} have successfully predicted barrier heights and rate constants of diverse gas phase reactions only based on readily available 2D information (e.g. SMILES strings) of reactants and products. However, such predictive models for liquid/solution phase reactions have been lightly investigated with limited applicability.²⁰

Solvents can have significant impacts on reaction rates and outcomes, and it is crucial to accurately predict these kinetic solvent effects. Recent research efforts have been devoted to employing ML (e.g. deep neural network) for free energy predictions of condensed phase reactions.^{15,18,19,21–26} Many of these studies^{18,19,21–23,26} combine the ML models with semi-empirical or lower-level QM/MM methods to obtain the energy predictions that match the accuracy of higher-level QM/MM methods. For example, Gómez-Flores et al.¹⁹ used a ML approach to predict the energy difference between the density functional tight-binding model and other higher level QM methods for a thiol-disulfide exchange reaction in water. In a study by Pan et al.,¹⁸ a ML model was trained to reproduce ab initio QM/MM potentials in free energy simulations for the aqueous Menshutkin reaction between ammonia and chloromethane. Farrar and Grayson²⁶ employed ML models to predict DFT-quality activation barriers for various nitro-Michael addition reactions in toluene based on the features generated from semi-empirical methods. These approaches, however, require semi-empirical QM/MM steps that are less suitable for instantaneous, automatic rate predictions. Furthermore, their models are limited to a single solvent and need the 3D coordinates or QM features of reactants and transition states as inputs, which are not readily available.

The ML models by Jorner et al.²⁴ and by Heid and Green¹⁵ are the few cases that can predict reaction properties in multiple solvents only based on the 2D structural information of molecules. Jorner et al.²⁴ employed a Gaussian process regression model and compared several 2D structural features to predict the barrier height of 443 S_NAr reactions in different solvents. In their work, the best accuracy was reached by adopting the BERT²⁷ reaction fingerprint. Heid and Green,¹⁵ on the other hand, used the condensed graph of reaction (CGR) as an input reaction representation for a graph convolutional neural network (GCNN). They applied the CGR GCNN model to the same S_NAr data set and were able to achieve better

barrier height predictions compared to the other models that used the BERT fingerprint or different reaction representations. While these models can provide fast kinetic estimations for solution-phase reactions at a low computational cost, only one reaction family was considered with a relatively small training set. A larger data set that contains more diverse types of reactions and solvents is needed in order to train a more generalized model for kinetic solvent effect predictions. Moreover, both models used fixed descriptors to represent solvents, but prior studies^{15,28,29} revealed that the learned molecular representations based on a graph convolutional approach outperform fixed molecular descriptors in many property prediction tasks.

In this study, we present a ML model that can predict kinetic solvent effects for a wide range of neutral reactions and solvents only based on atom-mapped reaction SMILES and solvent SMILES strings. More precisely, the model predicts the solvation free energy and solvation enthalpy of activation ($\Delta\Delta G_{\text{solv}}^\ddagger$, $\Delta\Delta H_{\text{solv}}^\ddagger$) for a reaction-solvent pair, which can be used to estimate a relative rate constant between a solution phase and a gas phase reaction or between the reaction in different solvents. Our model adopts a CGR GCNN architecture with separate GCNN layers for solvent molecular encoding. A large, diverse set of training data containing over 28,000 reactions and 295 solvents is generated in this work by performing ab initio COSMO-RS³⁰ calculations. The performance of the model on unseen reactions is rigorously assessed by comparing the ML predictions with both COSMO-RS calculations and experimental data. A transfer learning approach and various additional features are explored to further improve the model. Our ML model can provide accurate predictions of relative rate constants, and together with the existing predictive models or databases for gas phase rate constants (e.g. RMG database¹²), it can provide the estimates of absolute rate constants for many different liquid phase reactions.

2 Background on the prediction targets

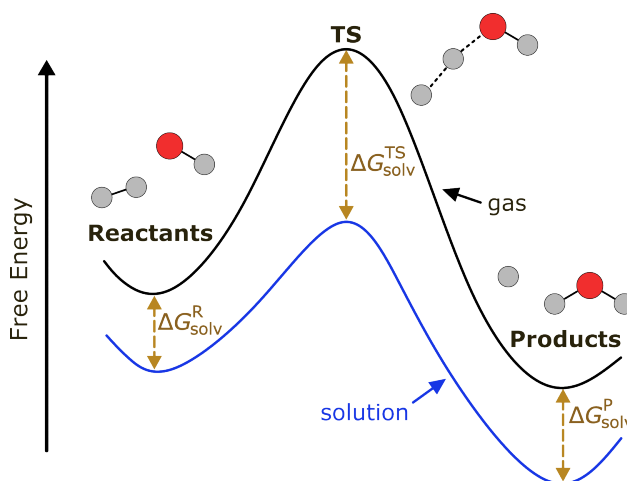


Figure 1: Potential energy diagram of a reaction in a gas phase and a solution phase.

Our ML model aims to predict the solvation free energy and solvation enthalpy of activation ($\Delta\Delta G_{\text{solv}}^\ddagger$, $\Delta\Delta H_{\text{solv}}^\ddagger$) at 298 K for a reaction in a solvent. Solvation free energy (ΔG_{solv}) and solvation enthalpy (ΔH_{solv}) are the changes in Gibbs free energy and enthalpy when a molecule is transferred from an ideal gas to a solvent at a fixed condition. The $\Delta\Delta G_{\text{solv}}^\ddagger$ and $\Delta\Delta H_{\text{solv}}^\ddagger$ of a reaction-solvent pair are defined as the solvation free energy and solvation enthalpy differences between a transition state (TS) and reactant(s):

$$\Delta\Delta G_{\text{solv}}^\ddagger = \Delta G_{\text{solv}}^{\text{TS}} - \Delta G_{\text{solv}}^{\text{R}} \quad (1)$$

$$\Delta\Delta H_{\text{solv}}^\ddagger = \Delta H_{\text{solv}}^{\text{TS}} - \Delta H_{\text{solv}}^{\text{R}} \quad (2)$$

where $\Delta G_{\text{solv}}^{\text{TS}}$ and $\Delta G_{\text{solv}}^{\text{R}}$ represent the solvation free energies of a TS and a reactant, and $\Delta H_{\text{solv}}^{\text{TS}}$ and $\Delta H_{\text{solv}}^{\text{R}}$ represent the solvation enthalpies of a TS and a reactant, respectively. For a bimolecular reaction, $\Delta G_{\text{solv}}^{\text{R}}$ and $\Delta H_{\text{solv}}^{\text{R}}$ each correspond to the sum of the solvation free energies and solvation enthalpies of all reactants. The standard state of 1 M ideal gas and 1 M solution is used for solvation free energy and enthalpy in this work.

As depicted in Figure 1, a solvent medium can affect the energies of reactants and a TS by different degrees, causing the activation free energy to shift when a reaction occurs in a solution (liquid) phase. The $\Delta\Delta G_{\text{solv}}^\ddagger$ of a reaction corresponds to the difference in the free energy of activation between a gas phase and a solution phase and is an important kinetic parameter for solution phase reactions. For example, $\Delta\Delta G_{\text{solv}}^\ddagger$ can be directly used to estimate the ratio of a gas phase rate constant (k_{gas}) to a liquid phase rate constant (k_{liq}) as follows:³¹

$$\frac{k_{\text{liq}}}{k_{\text{gas}}} = \exp\left(\frac{-\Delta\Delta G_{\text{solv}}^\ddagger}{RT}\right) \quad (3)$$

where R is the universal gas constant and T is a temperature. It can be also used to calculate the relative rate constant between two solvents:

$$k_{\text{rel}} = \frac{k_{\text{liq}}^{\text{s1}}}{k_{\text{liq}}^{\text{s2}}} = \exp\left(-\frac{\Delta\Delta G_{\text{solv,s1}}^\ddagger - \Delta\Delta G_{\text{solv,s2}}^\ddagger}{RT}\right) \quad (4)$$

where $k_{\text{liq}}^{\text{s1}}$ and $k_{\text{liq}}^{\text{s2}}$ are the rate constants of a reaction in a solvent 1 and in a solvent 2, respectively, and $\Delta\Delta G_{\text{solv,s1}}^\ddagger$ and $\Delta\Delta G_{\text{solv,s2}}^\ddagger$ are the corresponding solvation energies of activation for the reaction in each solvent.

Our model predicts $\Delta\Delta H_{\text{solv}}^\ddagger$ in addition to $\Delta\Delta G_{\text{solv}}^\ddagger$ at 298 K to account for the temperature dependence of $\Delta\Delta G_{\text{solv}}^\ddagger$. The $\Delta\Delta G_{\text{solv}}^\ddagger$ at a different temperature can be linearly extrapolated using the two model outputs at 298 K as follows:

$$\Delta\Delta G_{\text{solv}}^\ddagger(T) \approx \Delta\Delta H_{\text{solv}}^\ddagger(298\text{K}) - T \left(\frac{\Delta\Delta H_{\text{solv}}^\ddagger(298\text{K}) - \Delta\Delta G_{\text{solv}}^\ddagger(298\text{K})}{298\text{K}} \right) \quad (5)$$

The linear approximation is found to be generally valid for a moderate temperature range (250 K - 400 K),³² but the error is expected to increase as the temperature significantly deviates from 298 K.

3 Methods

3.1 Data generation

Table 1: Summary of the data sets used in this study. The number of reactions in the pre-training and fine-tuning sets include both forward and reverse directions. “N data chosen” represents the number of data sampled from the total data to construct the training, validation, and test sets.

Data set	N data total	N data chosen	N reactions	N solvents	Data type & reference
Pre-training set	7,796,583	500,000 (6.4 %)	26,448	295	In-house COSMO-RS calculations based on the optimized geometries from Grambow et al. ³³⁻³⁵
Fine-tuning set	542,833	46,122 (8.5 %)	1870	295	In-house COSMO-RS calculations based on the optimized geometries from Harms et al. ³⁶
Experimental test set	165	165	15	49	Experimental relative rate constants from Chung and Green ³¹

Table 1 shows the summary of the data sets used in this work. A total of three data sets are prepared: (1) a pre-training set containing the reactions from Grambow et al.,³³⁻³⁵ (2) a fine-tuning set containing the reactions from Harms et al.,³⁶ and (3) an experimental test set from our prior study.³¹ The data sets include diverse range of neutral closed-shell and free radical reactions and nonionic solvents. For both pre-training and fine-tuning sets, $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ are computed for each reaction-solvent pair with the COSMO-RS calculations based on the geometries obtained from Grambow et al. and Harms et al. The ML model is trained, validated, and tested on the computed data, and the experimental set is used as an additional test set for the final error assessment. The reaction and solvent information is stored as atom-mapped reaction SMILES and solvent SMILES in all data sets.

We separated the computed data into the pre-training and fine-tuning sets because the two data sets significantly differ in the types of reactions included and the level of theory used for geometry optimizations. The pre-training set is the largest, but the majority of its reactions are uncommon reactions with high gas phase barrier heights (e.g. $E_a > 50$ kcal/mol), and it does not contain any reactions that are bimolecular in both forward and reverse directions. In contrast, the fine-tuning set is smaller but contains more common reactions. To leverage the different types of data, we employ a transfer learning approach in which the model is

first pre-trained on the reactions from Grambow et al. and subsequently fine-tuned on the reactions from Harms et al. Details on each data set and the computational method are described below, and all data sets are provided as a part of the Supporting Information.

3.1.1 Computational method

The pre-training and fine-tuning data sets are generated by performing COSMO-RS calculations at the BP86/def2-TZVPD^{37–39} level of theory with fine grid cavity,⁴⁰ which is commonly known as a BP-TZVPD-FINE level. The COSMO-RS is a hybrid solvation model that uses quantum chemistry and statistical thermodynamics to compute the chemical potential of a compound in a solvent.^{30,41,42} We have previously demonstrated that the COSMO-RS method can provide accurate predictions of $\Delta\Delta G_{\text{solv}}^{\ddagger}$ for various neutral closed-shell and free radical reactions in different solvents with a mean absolute error of around 0.45 kcal/mol.³¹

The computational workflow used in this work follows that employed in our earlier study.³¹ Single-point energy calculations are performed at the BP-TZVPD-FINE level of theory in a COSMO phase and in a gas phase with TURBOMOLE 7.5^{43,44} for reactants, products, and TSs based on the optimized gas phase geometries obtained from Grambow et al. and Harms et al.; this step generates screening charge densities and energies that are needed for the COSMO-RS calculations. Then, the ΔG_{solv} and ΔH_{solv} of reactants, products, and TSs are computed in 295 common solvents at 298 K with COSMOtherm (release 2021)⁴⁵ based on the COSMO-RS theory, and the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ of each reaction are subsequently calculated in 295 solvents at 298 K using equations 1 and 2. Because COSMOtherm does not directly output solvation enthalpy, $\Delta\Delta H_{\text{solv}}^{\ddagger}$ is obtained by first computing $\Delta\Delta G_{\text{solv}}^{\ddagger}$ at 297, 298, and 299 K, estimating the temperature gradient at 298 K, and then using the definition $\Delta H = \Delta G - T \frac{d\Delta G}{dT}$. The screening charge densities and energies of the 295 solvents are acquired from the COSMObase database.⁴⁶ The reactions from Grambow et al. and Harms et al. were each optimized in gas phase at the ω B97XD3/def2-TZVP^{47,48} and M06-2X/cc-pVTZ^{49,50} levels of theory in their original work. Although these levels of theory are different from the level used for the COSMO-RS calculations, our prior work³¹ demonstrated that the accurate COSMO-RS calculations can be made with the gas phase geometries that are optimized at different levels of theory, which justifies the current computational workflow.

A total of 7,814,610 and 614,780 COSMO-RS calculations were completed successfully for the pre-training and fine-tuning sets, respectively. The results were then further cleaned by only including the reaction-solvent pairs that successfully ran for both forward and reverse directions of the reaction. While most of the computed $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ values were within ± 10 kcal/mol, a small fraction of the data had unreasonably large values for neutral reactions. For instance, a $\Delta\Delta G_{\text{solv}}^{\ddagger}$ of ± 40 kcal/mol corresponds to around 29 orders of magnitude increase/decrease in a liquid phase rate constant compared to a gas phase rate constant (see eq 3). We suspected that these are likely due to the COSMO-RS calculation errors as the COSMO-RS method may not have been parameterized well for certain reactions

and geometries. Therefore, we filtered out 241 reaction-solvent pairs from the pre-training set that had $|\Delta\Delta G_{\text{solv}}^\ddagger| > 40$ kcal/mol or $|\Delta\Delta H_{\text{solv}}^\ddagger| > 56$ kcal/mol (14 standard deviations away from means). Higher quality data are usually expected for the fine-tuning set. Thus, more strict cutoff values of $|\Delta\Delta G_{\text{solv}}^\ddagger| > 10$ kcal/mol and $|\Delta\Delta H_{\text{solv}}^\ddagger| > 18$ kcal/mol are applied to the fine-tuning set to remove potentially erroneous data.

3.1.2 Pre-training set

The final pre-training set contains a total of 7,796,583 reaction-solvent pairs with 26,448 unique reactions and 295 solvents. Both forward and reverse reactions are included in the data set to augment the data. As mentioned earlier, the geometry optimizations were done at the ω B97XD3/def2-TZVP level of theory for these reactions in the original work by Grambow et al.³³⁻³⁵ The histograms and statistics of the data set are provided in Supporting Information Figure S1. The computed $\Delta\Delta G_{\text{solv}}^\ddagger$ and $\Delta\Delta H_{\text{solv}}^\ddagger$ have nearly normal distributions with high peaks at zero. The $\Delta\Delta G_{\text{solv}}^\ddagger$ and $\Delta\Delta H_{\text{solv}}^\ddagger$ have absolute mean values of 1.81 and 2.58 kcal/mol, respectively, with standard deviations of 2.76 and 3.92 kcal/mol. The reactions contain diverse types of neutral closed-shell and free radical reactions that involve H, C, N, and O atoms and have at most 7 heavy atoms. Due to errors from Open Babel⁵¹ when perceiving connectivity, a small set of the original reaction SMILES from the Grambow et al. had incorrect bond orders and formal charges, and therefore, the corrected atom-mapped SMILES from Spiekermann et al.⁵² are used for our pre-training set.

The entire data set has nearly 7.8 million data points. However, it is unlikely that every reaction-solvent pair is needed since the total number of unique reactions and solvents would remain fixed even if the number of reaction-solvent pairs increases. To investigate the effect of the data size on the model performance, we prepared 8 different data sets containing 10k, 50k, 75k, 100k, 250k, 500k, 750k, and 1M data points. These data are sampled in a semi-random manner such that all reactions and solvents appear in the data sets at least once, except the 10k set which has fewer data than the total number of reactions. From the results, we determined 500k to be the optimal data set size for the model as explained further in the results section.

3.1.3 Fine-tuning set

The fine-tuning data set has 542,833 reaction-solvent pairs with 1870 unique reactions including both forward and reverse directions and 295 solvents. The geometry optimizations were performed at the M06-2X/cc-pVTZ level of theory in the original work by Harms et al.³⁶ for these reactions. The data set contains three specific reaction types: bimolecular hydrogen abstraction (H-abstraction), unimolecular hydrogen migration (H-migration), and radical addition to a multiple bond (R-addition). These are neutral, free radical reactions that are ubiquitous in both gas and liquid phase systems. The reaction templates are illus-

trated in the Supporting Information Figure S2. In total, there are 1402 H-abstraction, 146 H-migration, and 322 R-addition reactions. The reactions involve maximum 10 heavy atoms and include H, C, and O atoms. The histograms and statistics of the data are presented in Supporting Information Figure S3. The $\Delta\Delta G_{\text{solv}}^\ddagger$ and $\Delta\Delta H_{\text{solv}}^\ddagger$ have absolute average values of 1.40 and 2.32 kcal/mol, respectively, with standard deviations of 1.91 and 3.17 kcal/mol.

Similar to the pre-training set, only a subset of the fine-tuning data was chosen for the model. We sampled around 25 solvents per each reaction semi-randomly with more weights on polar solvents in order to include more data with stronger solvent effects. A total of 46,122 data points were selected, and all reactions and solvents appear in the chosen set at least once. The detailed data sampling method is explained in Supporting Information Section S2.

3.1.4 Experimental test set

The experimental data set³¹ consists of 165 relative rate constants (k_{rel}) for 15 neutral reactions and 49 solvents from 273 K to 392 K. None of the experimental reactions appear in the pre-training and fine-tuning sets, and therefore these data serve as a reaction split test set for the final model evaluation. The reactions are depicted in Supporting Information Table S1, and they include 2 β -scission, 5 H-abstraction, 3 Diels-Alder, and 5 other types of reactions. The reactions involve H, C, N, O, and S atoms and have up to 22 heavy atoms, which are much larger than the reactions found in the pre-training and fine-tuning sets. The errors on the experimental test set are computed in both $\log_{10}(k_{\text{rel}})$ and $\Delta G_{\text{rel}}^\ddagger$ units, where the $\Delta G_{\text{rel}}^\ddagger$ error is calculated as follows:

$$\Delta G_{\text{rel}}^\ddagger \text{ error} = -RT (\ln(k_{\text{rel,expt}}) - \ln(k_{\text{rel,calc}})) \quad (6)$$

3.2 Data splits

The pre-training set is split into a 90% training/validation and a 10% test set using reaction and solvent splits to evaluate the model's predictive performance on unseen reactions and unseen solvents. For our test splits, 5% reactions and 5% solvents are randomly selected, and all reaction-solvent pairs that include the chosen reactions or chosen solvents are added to the test set and excluded from the training/validation set. Both forward and reverse directions of the selected reactions are included in the test set to prevent data leakage; this procedure is very crucial for evaluating the true performance of a model on unseen reactions since the model can gain the information on the reaction and TS from its own reverse reaction. The importance of the proper reaction split is addressed in the recent studies by Heid and Green¹⁵ and Spiekermann et al.¹⁶ The remaining 90% data are randomly split into a 80% training and 20% validation set. The validation set is used for early stopping to determine the epoch that gives the lowest validation error and prevents over-fitting. Five folds are used to prepare five different training, validation, and test sets from the pre-training set.

The fine-tuning set is randomly split into a 80 % training and a 20 % validation set using one fold. We did not construct a separate test set from the fine-tuning set. Instead, the experimental set is used as a final reaction-split test set for the fine-tuned model. We ensured that none of the experimental reactions (both forward and reverse) appear in the pre-training and fine-tuning sets.

3.3 Machine learning model architecture




Figure 2: Schematic of a machine learning model architecture. The model takes an atom-mapped reaction SMILES and a solvent SMILES as inputs.

The schematic of the ML architecture is depicted in Figure 2. Our model takes the atom-mapped reaction SMILES and solvent SMILES as inputs and predicts $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ at 298 K. The model is constructed using Chemprop,²⁸ an open-source software that uses a directed message passing neural network (D-MPNN) for chemical and reaction property predictions. A D-MPNN is a type of GCNN that converts atom and bond features into a latent representation of a molecule through bond-level message passing. To encode a reaction, we adopt the established CGR representation^{15,53,54} as it has shown to outperform other representations for various reaction property predictions. The CGR is a superposition of the reactant and product graphs, which mimics the 2D-structure of the TS. In our model, the CGR representation is constructed from the atom-mapped reactants and products and passed into a D-MPNN to give a reaction embedding. A separate D-MPNN is employed to convert a solvent graph into a solvent molecular embedding. The learned reaction and solvent representations are then concatenated together and passed into a feed forward neural network (FNN) to predict the regression targets. The initial atom and bond features are generated using RDKit⁵⁵ within Chemprop and include several features such as atom type, bond type, formal charge, chirality, and etc. The hyperparameters are optimized via 47 iterations of

Bayesian optimization with the Hyperopt package.⁵⁶ Only the training/validation set of the pre-training data is used for the hyperparameter optimization to prevent data leakage. The full list of atom and bond features and the optimized hyperparameters can be found in Supporting Information Tables S2 and S3.

As mentioned in the earlier section, a transfer learning approach is used to first train the model on the pre-training set and subsequently fine-tune the model on the fine-tuning set with fewer epochs. Within each of the 5 folds of the pre-training set, an ensemble of 5 different models are generated by using different random initialization of model parameters. This results in a total of 25 individual models from the 5-fold pre-training set. The optimized parameters from the 25 pre-trained models are then used to initialize the 25 fine-tuned models that are trained on the fine-tuning set. When evaluating the reaction- and solvent-split errors on the pre-training test set, average predictions from the ensemble of 5 models are employed to compute the error in each fold. When evaluating the final error on the experimental set, an ensemble of all 25 fine-tuned models are used to make the average predictions. We ensured that no models are trained or validated on the tested reactions or solvents.

3.4 Additional features

Table 2: List of additional features investigated.

Name	Type	Description	Ref
RP-solv	Molecular	Solvation energy and enthalpy of reactants and products at 298 K ($\Delta G_{\text{solv}}^{\text{R}}$, $\Delta G_{\text{solv}}^{\text{P}}$, $\Delta H_{\text{solv}}^{\text{R}}$, $\Delta H_{\text{solv}}^{\text{P}}$)	This work
S-abr	Molecular	Abraham parameters of solvents	57
RDKit-mol	Molecular	2D molecular features generated from RDKit	55
Addit-atom	Atomic	Additional atom features generated from RDKit	55,58
QM-desc	Atomic, bond	QM atom and bond descriptors	59

Five additional sets of features that are listed in Table 2 are explored to improve the model performance. The RP-solv features represent the solvation free energies and solvation enthalpies of reactants and products that are calculated in this work. We also tested the Abraham parameters of solvents (S-abr) that are obtained from the ML model by Chung et al.⁵⁷ The Abraham parameters consist of five descriptors that can describe various solute/solvent interactions.⁶⁰ The RDKit-mol represents the 2D molecular features generated from RDKit. There are 200 2D molecular features available within RDKit, and 20 and 15 features were selected for reactants/products and solvents, respectively, based on the variance threshold and random forest methods imported from scikit-learn.⁶¹ Six additional RDKit atomic features (Addit-atom) that are adopted from Vermeire and Green⁵⁸ are also tested as they have shown to improve solvation free energy predictions. These include the number of radical electrons, ring size, number of lone electron pairs, H-bond donating and accepting characters, and electronegativity. Lastly, the QM atomic and bond descriptors (QM-desc)

obtained from the ML model by Guan et al.⁵⁹ are explored. The QM-desc contains 4 atomic descriptors (Hirshfeld partial charge, two Fukui indices, NMR shielding constants) and 2 bond descriptors (bond lengths, bond orders) that were shown to improve reaction property predictions.^{59,62}

The performances of the additional features are compared using the pre-training test set. Within the ML model, additional molecular features are concatenated with the reaction and solvent embeddings and fed into the FFN layer to make the predictions. Additional atom and bond features are concatenated with the initial atom and bond features prior to the CGR/D-MPNN layers. Note that nearly all features can be calculated instantly or predicted by existing ML models. The only exception is the RP-solv features which are computed with the COSMO-RS method. Yet, several ML models are available for predicting solvation energy and enthalpy of closed-shell compounds,^{57,58,63–66} and the RP-solv features can be therefore estimated with the ML models if fast approximations are needed. We did not consider the 3D structures of the reactants and products as additional inputs in our study as they are usually not readily available and prone to calculation noise and error. Furthermore, Spiekermann et al.¹⁶ showed that the 2D D-MPNN model outperformed the 3D ML model for gas phase barrier height predictions on the Grambow et al.’s reactions. Since the same data set and similar model architecture are used in our study, we expect the result to be similar and hence do not consider the 3D ML model in this work.

4 Results and discussion

4.1 Data set size and additional features

The effects of the data set size and additional features are investigated using the pre-training set prior to fine-tuning any models. The resulting test root-mean-square errors (RMSE) on the reaction and solvent splits are presented in Figure 3. The reaction and solvent splits each test the model’s performance on unseen reactions in seen solvents and on seen reactions in unseen solvents. As previously explained, the pre-training set of around 7.8M data is divided into smaller subsets to identify the optimal data set size that can balance accuracy and training time. From Figure 3a, it can be seen that the test error initially decreases with an increasing data set size and plateaus out from 500k for the reaction split. For the solvent split, the error continues to decrease at a higher data set size, but the change in the error is very small beyond 500k. Therefore, 500k is chosen as a final data set size for the pre-training set. This result was expected since only the number of reaction-solvent pairs increases with the increasing data set size whereas the number of unique reactions and solvents remains constant. The information gain from more reaction-solvent pairs is likely to saturate after the model sees enough data on each reaction and solvent, causing the errors to level out.

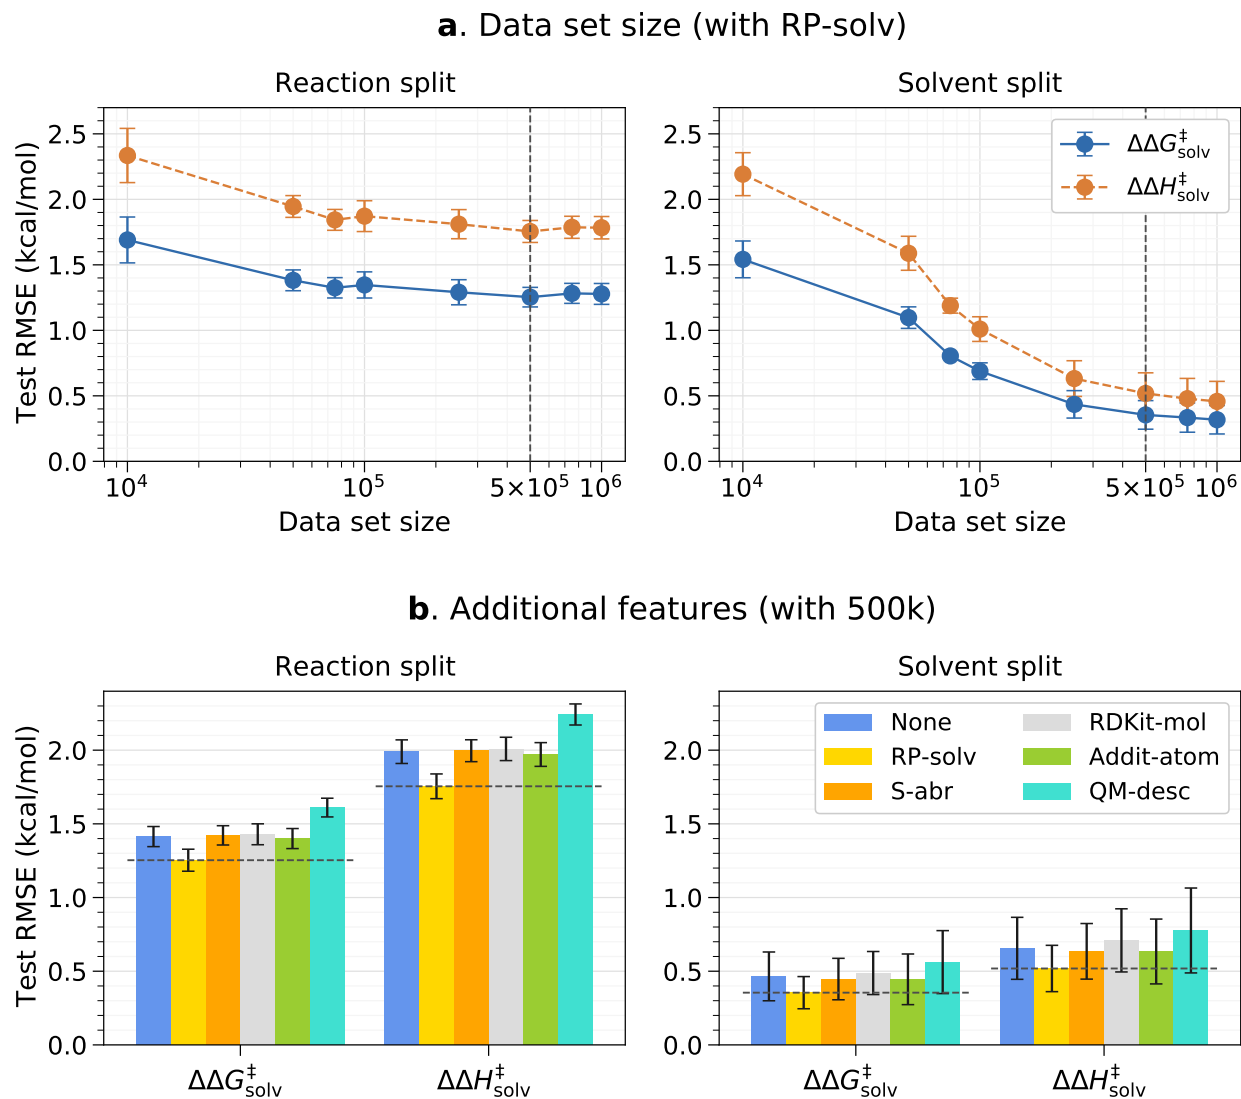


Figure 3: Pre-trained model results on the reaction split and solvent split test sets. (a) RMSE vs the data set size for the model trained with the RP-solv feature. (b) RMSE error of different additional features for the model trained with the 500k data set. The chosen data set size and feature are marked with dashed vertical and horizontal lines, respectively. The error bars indicate the standard deviation between five folds.

Figure 3b shows the results of the additional features tested with the 500k data set. In all cases, only the RP-solv feature improves the errors and other features do not have noticeable impacts on the model performance. The reactions tend to accelerate/decelerate in a polar solvent when the products are more/less strongly solvated than the reactants.^{67,68} The RP-solv feature, which consists of solvation energy and solvation enthalpy of reactants and products, describe how strongly the reactants and products are solvated in a solvent and therefore helps the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ predictions for our model. The results also show that the QM-desc feature worsens the model performance, which is consistent with results from Spiekermann et al.¹⁶ that show the descriptors offer no improvement. The QM-desc

feature was predicted by the ML model from Guan et al.⁵⁹ which was trained on three classes of substitution reactions. Our data set includes more diverse types of reactions, and therefore, it is likely that their model is not suitable for our reaction targets. Moreover, it is possible that the chosen QM descriptors are not related to our targets and act as noise to the model.

4.2 Comparison of the pre-trained and fine-tuned models

Table 3: Test errors of different models and additional (addit.) features. The best combination of the model and additional feature is marked in bold. For the pre-training set, the errors on the reaction split are available for all models while the errors on the solvent split are only available for the pre-trained model and shown in parentheses. The standard deviations are calculated between the 5 folds for the pre-trained set. All test errors are reported in kcal/mol.

Model	Addit. feature	Pre-training set:				Experimental set	
		reaction split		(solvent split)		$\Delta G_{\text{rel}}^{\ddagger}$	
		$\Delta\Delta G_{\text{solv}}^{\ddagger}$		$\Delta\Delta H_{\text{solv}}^{\ddagger}$		MAE	RMSE
		MAE	RMSE	MAE	RMSE		
Pre-trained	None	0.79 ± 0.03 (0.20 ± 0.05)	1.41 ± 0.07 (0.46 ± 0.17)	1.13 ± 0.04 (0.30 ± 0.07)	1.99 ± 0.08 (0.66 ± 0.21)	0.90	1.18
Pre-trained	RP-solv	0.68 ± 0.02 (0.16 ± 0.04)	1.25 ± 0.07 (0.35 ± 0.11)	0.97 ± 0.03 (0.24 ± 0.06)	1.75 ± 0.08 (0.52 ± 0.16)	0.75	0.98
Fine-tuned	None	0.71 ± 0.14 (-)	1.16 ± 0.24 (-)	1.03 ± 0.18 (-)	1.63 ± 0.33 (-)	0.60	0.80
Fine-tuned	RP-solv	0.82 ± 0.08 (-)	1.29 ± 0.18 (-)	1.40 ± 0.09 (-)	2.06 ± 0.21 (-)	0.71	0.99

Table 3 summarizes the performances of the pre-trained and fine-tuned models with and without the RP-solv feature. The MAE (mean absolute error) and RMSE are reported in kcal/mol, and the standard deviations are calculated between the 5 folds for the pre-training set. For the experimental test set, the k_{rel} error is converted to a Gibbs free energy unit ($\Delta G_{\text{rel}}^{\ddagger}$ or $RT(\ln k_{\text{rel}})$) using eq 6 to allow easy comparison.

The fine-tuned model with no additional feature achieves overall the best performance on unseen reactions for both pre-training and experimental test sets and is therefore selected as the optimal model. The model has the $RT(\ln k_{\text{rel}})$ MAE/RMSE of 0.60/0.80 kcal/mol on the experimental set and has the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ MAE/RMSE of 0.71/1.16 and 1.03/1.63 kcal/mol, respectively, on the pre-training set reaction split. The model has higher $\Delta\Delta H_{\text{solv}}^{\ddagger}$ errors than $\Delta\Delta G_{\text{solv}}^{\ddagger}$ in all cases as the COSMO-RS method, which was used to generate the training data, has higher calculation errors for $\Delta\Delta H_{\text{solv}}^{\ddagger}$.⁵⁷ Furthermore, $\Delta\Delta H_{\text{solv}}^{\ddagger}$ generally has a larger magnitude than $\Delta\Delta G_{\text{solv}}^{\ddagger}$, which leads to larger absolute errors. The results also show that the pre-trained model has much lower errors on the solvent split than the reaction

split. The model is able to provide very accurate predictions on unseen solvents with 295 training solvents whereas it has much higher errors on unseen reactions even with 26,448 training reactions. We believe this is because the chemical space of viable solvents is not as nearly big as that of reactions. Furthermore, the reaction split is a more challenging task since the model has to infer the TS information from the reactants and products. The fine-tuned model was not separately tested on the solvent split as it was trained on all solvents that are found in the fine-tuning set. Since the major limitation is on the reaction split, we expect the fine-tuned model to have a similarly low error on unseen solvents.

Contrary to the earlier results on the pre-trained model, it is found that the RP-solv feature has adverse effect on the fine-tuned model. It is likely that the model was able to learn more comprehensive reaction embeddings from the fine-tuning data set and no longer needs additional information on reactants and products. As a result, the RP-feature serves as an unnecessary input and hinders the performance of the new model. It is surprising that fine-tuning the model with a relatively smaller data set can lead to a completely different outcome for the feature selection. Compared to the pre-trained model, the fine-tuned model not only has better performance but also benefits from a reduced computational cost as it no longer requires the RP-solv feature that needs to be calculated.

Table 3 shows that the best fine-tuned model achieves around 0.18 kcal/mol lower error than the pre-trained model on the unseen experimental reactions. The performance gain is relatively big considering the small size of the fine-tuning set compared to the size of the pre-training set. The fine-tuning set contains bimolecular reactions that the pre-training set lacks and includes more common classes of reactions while the pre-training set largely contains uncommon reactions with high gas phase barrier heights ($E_a > 50$ kcal/mol).³³ Hence, even a relatively small number of fine-tuning data greatly enhances the model's performance on the experimental set, which mostly contains low-barrier reactions and several bimolecular reactions. A similar result was observed in the work by Spiekermann et al.¹⁶ where a model that was initially pre-trained with lower accuracy DFT data showed substantial improvement on barrier height predictions after fine-tuning with a small number of higher accuracy CCSD(T)-F12 data. Both their and our studies demonstrate that different types of data sets can be best leveraged via transfer learning when only a limited amount of higher quality or more relevant data is available. Transfer learning is particularly beneficial for our study since we could avoid mixing the two data sets that differ in the level of theory used for geometry optimizations and also put more emphasis on the data set that is considered to be more relevant to real liquid phase systems.

It is also worthwhile to note that the model has similar or slightly lower errors on the pre-training set reaction split after fine-tuning. Even though the pre-training and fine-tuning sets differ in the level of theory used for geometry optimizations and the types of reactions included, fine-tuning improves the model's performance on the pre-training test set as well. Our prior study³¹ demonstrated that the $\Delta\Delta G_{\text{solv}}^\ddagger$ calculations using the COSMO-




Figure 4: Model MAE vs the number of fine-tuning epochs. The model is trained with no additional feature. The error on the pre-training set is evaluated on the reaction split test set, and the error on the fine-tuning set is evaluated on the random split validation set.

RS method are not too sensitive to the level of theory used for geometry optimizations for the 15 experimental reactions tested. Similar conclusion can be deduced from the current result as the fine-tuning set, which is based on the M06-2X/cc-pVTZ geometries, still helps or does not exacerbate the model's predictions on the pre-training set, which is based on the ω B97XD3/def2-TZVP geometries.

However, the model can have very different outcomes depending on the number of fine-tuning epochs used. In this work, the ML model was trained up to the chosen number of maximum epochs, and the validation set was then used to determine at which epoch to stop the model training. We used the maximum epoch of 80 for pre-training and used the smaller maximum epoch of 10 for fine-tuning to prevent the pre-trained information from being completely overwritten by the three reaction families used in the fine-tuning set. Moreover, the error on the fine-tuning validation set plateaued out after 10 epochs (see Figure 4), and thus, the maximum fine-tuning epoch was selected as 10 in our work. To investigate the effect of fine-tuning epochs on the model performance, we trained additional models using different maximum epochs and compared the results. Figure 4 shows that the error on the pre-training set initially has a sharp drop as the model learns new reactions but gradually increases as the number of maximum fine-tuning epochs increases. The error on the experimental set also decreases at first but soon reaches a plateau at around 10 epochs. The result indicates that the chosen epoch maintains a good balance between retaining previously learnt knowledge and learning new data. At higher epochs, however, the model starts to lose prior knowledge without much added benefits as it becomes biased toward the three reaction families found

in the fine-tuning set. It is thus important to identify optimal epochs and hyperparameters for the fine-tuned model if one seeks to preserve the pre-trained information.

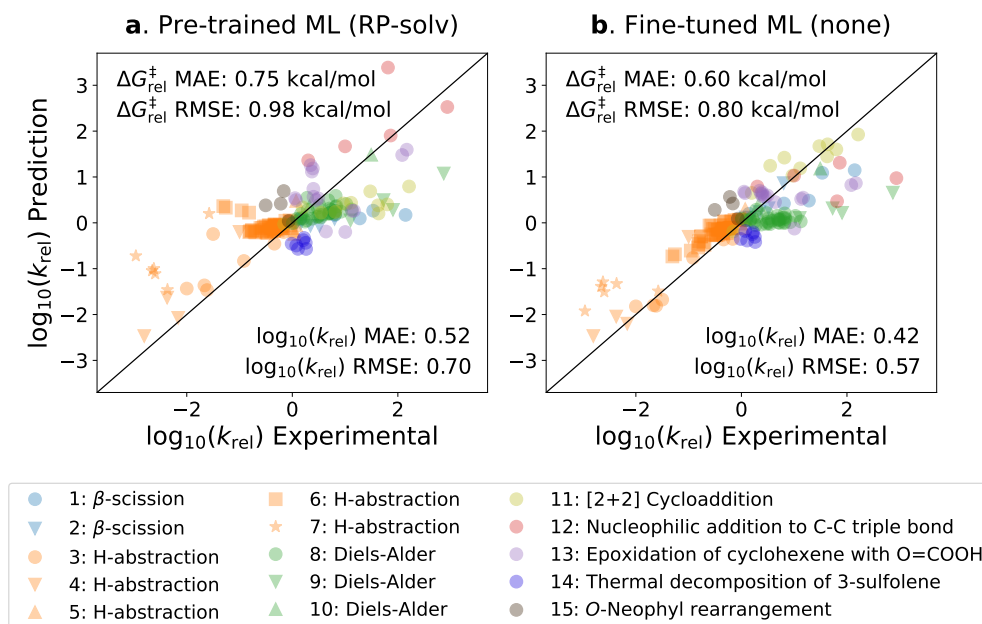


Figure 5: Parity plots of the predicted vs experimental k_{rel} values. (a) The pre-trained model with the RP-solv feature. (b) The fine-tuned model with no additional feature. The errors are reported in both $\log_{10}(k_{\text{rel}})$ and $\Delta G_{\text{rel}}^{\ddagger}$ units.

The parity plots of the best pre-trained model (with RP-solv) and fine-tuned model (no additional feature) are presented in Figure 5 for the experimental test set. It can be seen that the predictions on the H-abstraction reactions are substantially improved after fine-tuning. This was expected as the fine-tuning set primarily comprises H-abstraction reactions. We find that fine-tuning also improves the predictions on other reactions such as [2+2] cycloaddition and β -scission. The model, on the contrary, has relatively poor performance on Reactions 9, 12, and 13 (see Supporting Information Table S1 for details on the reactions). However, considering that the reactants and TSs found in the experimental set are around twice as large as those of the training sets, the model's extrapolation capability is quite satisfactory. Further improvement of the model is expected if more diverse and larger reactions become available as training data. In this work, the model was trained with the calculated data, which inherently have some errors. Our previous study showed that the COSMO-RS method has a $\Delta G_{\text{rel}}^{\ddagger}$ MAE/RMSE of around 0.4 kcal/mol (0.28 in $\log_{10} k_{\text{rel}}$ units) on these reactions. This suggests that only a part of the prediction errors are due to the model performance, and the rest are from the intrinsic calculation errors in the training data.

The parity plots and error histograms of the models on the pre-training set reaction split are provided in Figure 6 for the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ predictions. The corresponding plots for the $\Delta\Delta H_{\text{solv}}^{\ddagger}$ predictions and for the solvent split are presented in Supporting Information Figure S4 and

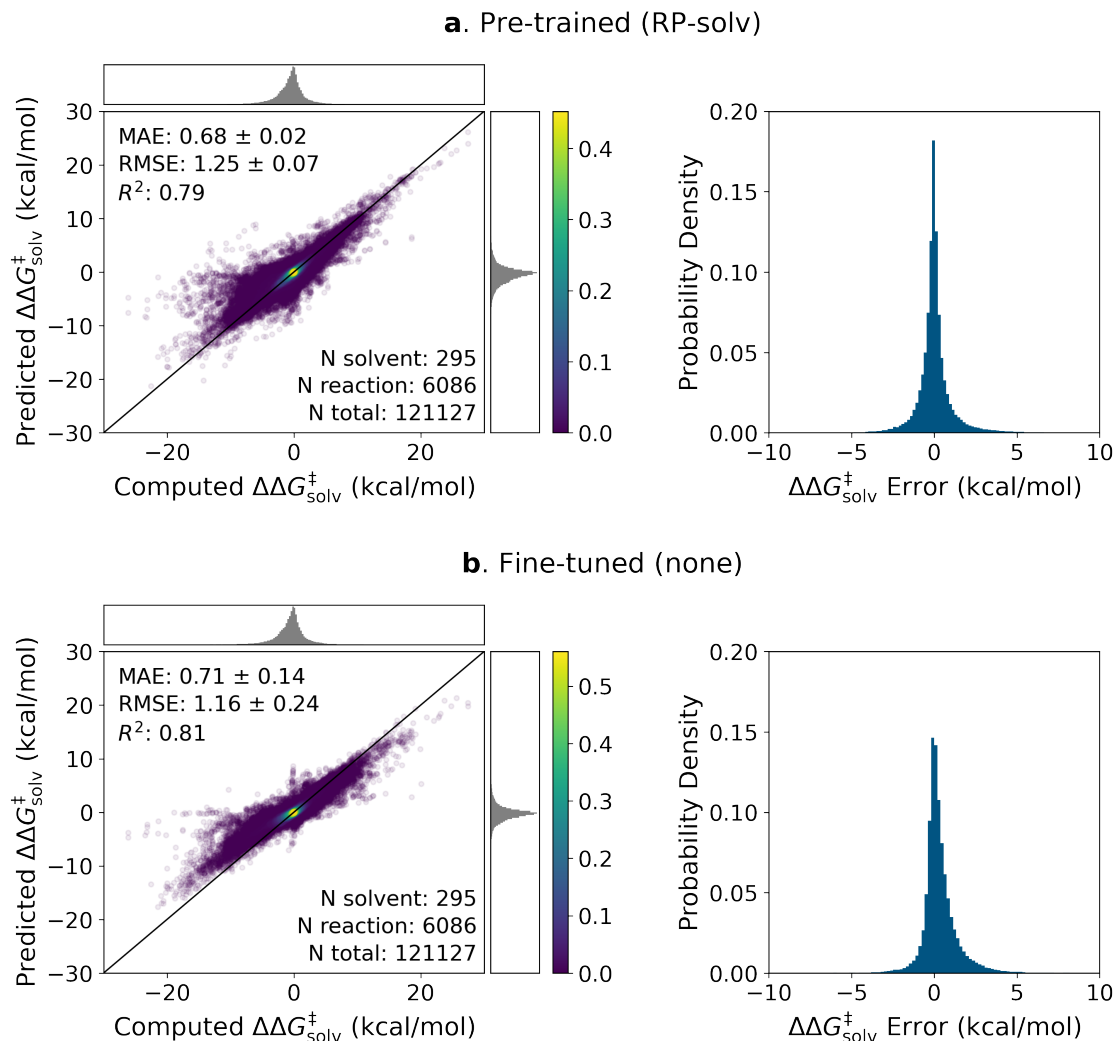


Figure 6: Parity plots and histograms of the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ errors on the 5-fold pre-training set reaction split. (a) The pre-trained model with the RP-solv feature. (b) The fine-tuned model with no additional feature. The MAE and RMSE are in kcal/mol. The numbers of reactions, solvents, total data points found in the test set are provided. The top and right subfigures on the parity plots show the distribution of computed and predicted values, and the colorbars display the scale of the 2D kernel density estimate plots.

Figure S5. The test errors are centered around zero for both models, and the majority of the errors fall within ± 2 kcal/mol. Yet, higher errors are observed in regions where the data are sparse. The pre-trained model predicts nearly zero $\Delta\Delta G_{\text{solv}}^{\ddagger}$ values for many reaction-solvent pairs whose computed $\Delta\Delta G_{\text{solv}}^{\ddagger}$ values are highly negative. Such trend is less pronounced in the fine-tuned model, but the model still tends to underpredict the magnitude of the $\Delta\Delta G_{\text{solv}}^{\ddagger}$ values when the computed $\Delta\Delta G_{\text{solv}}^{\ddagger}$ have large positive or negative values. It should be noted, though, that the model is compared with the computed values and not with the true values. The $\Delta\Delta G_{\text{solv}}^{\ddagger}$ of ± 10 kcal/mol corresponds to around 7 orders of magnitude difference between the liquid phase and gas phase rate constants at room temperature. Such

large solvent effects are very rare for neutral reactions, and thus, it is possible that the COSMO-RS method overpredicted the magnitude of the target values on some of these extreme data points.

We further examined the outliers of the fine-tuned model on the pre-training set reaction split. The top 20 reactions with the highest test errors are given in Supporting Information Table S4 along with their gas phase barrier heights obtained from Grambow et al.³³⁻³⁵ It is found that the majority of the outliers are unusual reactions such as those forming biradical products, involving TSs with high ring strain, and with high barrier heights ($E_a > 85$ kcal/mol). These reactions are unlikely to occur in real condensed phase systems, and therefore we expect the model to have lower errors on more feasible reactions.

5 Conclusions

We developed a machine learning model that can provide fast and accurate predictions of kinetic solvent effects for a wide range of neutral reactions and solvents. A large set of training data were generated using the COSMO-RS method for over 28,000 reactions and 295 solvents. The performance of the model was evaluated with both calculated and experimental data using rigorous data splits. The model achieves the MAEs of 0.71 and 1.03 kcal/mol on unseen reactions for the prediction of solvation free energy and solvation enthalpy of activation ($\Delta\Delta G_{\text{solv}}^\ddagger$, $\Delta\Delta H_{\text{solv}}^\ddagger$), respectively, relative to the COSMO-RS calculations. The model is shown to provide reliable predictions of relative rate constants when tested on the experimental set that contains unseen reactions with much bigger molecules than those found in the training set. We also demonstrate that different types of data sets can be effectively used via a transfer learning approach to refine the predictions.

The presented model can be used to estimate the relative rate constants between a gas phase and a liquid phase or between two solvents for a temperature range of around 250 K to 400 K. If a rate constant in one solvent or in a gas phase is known for a reaction, our model outputs can be used to estimate absolute rate constants in many different solvents for a given reaction. One of the advantages of the model is that it only needs the atom-mapped reaction SMILES and solvent SMILES as inputs, which are more amenable for automatic high-throughput predictions in comparison to requiring optimized 3D geometries as input. We anticipate the model to be particularly useful for the design of chemical processes and automatic construction of reaction mechanisms where fast estimations of kinetic parameters and solvent screenings are needed for a large number of reaction-solvent pairs.

6 Data and model availability

All data sets and the fine-tuned ML model can be found through Zenodo: <https://zenodo.org/record/8049537>. A sample script for making $\Delta\Delta G_{\text{solv}}^{\ddagger}$ and $\Delta\Delta H_{\text{solv}}^{\ddagger}$ predictions with the fine-tuned ML model can be found at https://github.com/yunsiechung/chemprop/tree/RxnSolvKSE_ML. The data sets and model are open access and distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license (<https://creativecommons.org/licenses/by/4.0/>).

Acknowledgement

The authors gratefully acknowledge the support of Eni S.p.A. for funding this research and the MIT SuperCloud Lincoln Laboratory Supercomputing Center⁶⁹ for providing HPC resources. The authors also thank Professor Richard West and his group for providing the reaction data set, Oscar Wu for his help on data compilation, and Kevin Spiekermann for suggestions that improved this manuscript.

References

- (1) Vermeire, F. H.; Aravindakshan, S. U.; Jocher, A.; Liu, M.; Chu, T.-C.; Hawtof, R. E.; Van de Vijver, R.; Prendergast, M. B.; Van Geem, K. M.; Green, W. H. Detailed Kinetic Modeling for the Pyrolysis of a Jet A Surrogate. *Energy Fuels* **2022**, *36*, 1304–1315.
- (2) Payne, A. M.; Spiekermann, K. A.; Green, W. H. Detailed Reaction Mechanism for 350–400 °C Pyrolysis of an Alkane, Aromatic, and Long-Chain Alkylaromatic Mixture. *Energy Fuels* **2022**, *36*, 1635–1646.
- (3) Edeleva, M.; Van Steenberge, P. H.; Sabbe, M. K.; D’hooge, D. R. Connecting gas-phase computational chemistry to condensed phase kinetic modeling: The state-of-the-art. *Polymers* **2021**, *13*, 1–39.
- (4) Chatelain, K.; Nicolle, A.; Ben Amara, A.; Catoire, L.; Starck, L. Wide Range Experimental and Kinetic Modeling Study of Chain Length Impact on n-Alkanes Autoxidation. *Energy Fuels* **2016**, *30*, 1294–1303.
- (5) Wu, H.; Grinberg Dana, A.; Ranasinghe, D. S.; Pickard, F. C.; Wood, G. P. F.; Zellesky, T.; Sluggett, G. W.; Mustakis, J.; Green, W. H. Kinetic Modeling of API Oxidation: (2) Imipramine Stress Testing. *Mol. Pharm.* **2022**, *19*, 1526–1539.
- (6) Yang, J.; Smith, M. C.; Prendergast, M. B.; Chu, T.-C.; Green, W. H. C₁₄H₁₀ polycyclic aromatic hydrocarbon formation by acetylene addition to naphthalenyl radicals observed. *Phys. Chem. Chem. Phys.* **2021**, *23*, 14325–14339.

- (7) Lu, T.; Law, C. K. Toward accommodating realistic fuel chemistry in large-scale computations. *Prog. Energy Combust. Sci.* **2009**, *35*, 192–215.
- (8) Sumathi, R.; Carstensen, H.-H.; Green, W. H. Reaction Rate Predictions Via Group Additivity. Part 3: Effect of Substituents with CH₂ as the Mediator. *J. Phys. Chem. A* **2002**, *106*, 5474–5489.
- (9) Saeys, M.; Reyniers, M.-F.; Marin, G. B.; Van Speybroeck, V.; Waroquier, M. Ab initio group contribution method for activation energies for radical additions. *AIChE J.* **2004**, *50*, 426–444.
- (10) Van de Vijver, R.; Sabbe, M. K.; Reyniers, M.-F.; Van Geem, K. M.; Marin, G. B. Ab initio derived group additivity model for intramolecular hydrogen abstraction reactions. *Phys. Chem. Chem. Phys.* **2018**, *20*, 10877–10894.
- (11) Gao, C. W.; Allen, J. W.; Green, W. H.; West, R. H. Reaction Mechanism Generator: Automatic construction of chemical kinetic mechanisms. *Comput. Phys. Commun.* **2016**, *203*, 212–225.
- (12) Johnson, M. S.; Dong, X.; Grinberg Dana, A.; Chung, Y.; Farina, D.; Gillis, R. J.; Liu, M.; Yee, N. W.; Blondal, K.; Mazeau, E.; Grambow, C. A.; Payne, A. M.; Spiekermann, K. A.; Pang, H.-W.; Goldsmith, C. F.; West, R. H.; Green, W. H. RMG Database for Chemical Property Prediction. *J. Chem. Inf. Model.* **2022**, *62*, 4906–4915.
- (13) Grambow, C. A.; Pattanaik, L.; Green, W. H. Deep Learning of Activation Energies. *J. Phys. Chem. Lett.* **2020**, *11*, 2992–2997.
- (14) Heinen, S.; von Rudorff, G. F.; von Lilienfeld, O. A. Toward the design of chemical reactions: Machine learning barriers of competing mechanisms in reactant space. *J. Chem. Phys.* **2021**, *155*, 064105.
- (15) Heid, E.; Green, W. H. Machine Learning of Reaction Properties via Learned Representations of the Condensed Graph of Reaction. *J. Chem. Inf. Model.* **2022**, *62*, 2101–2110.
- (16) Spiekermann, K. A.; Pattanaik, L.; Green, W. H. Fast Predictions of Reaction Barrier Heights: Toward Coupled-Cluster Accuracy. *J. Phys. Chem. A* **2022**, *126*, 3976–3986.
- (17) Johnson, M. S.; Green, W. H. A Machine Learning Based Approach to Reaction Rate Estimation. *ChemRxiv* **2022**, DOI: 10.26434/chemrxiv-2022-c98gc.
- (18) Pan, X.; Yang, J.; Van, R.; Epifanovsky, E.; Ho, J.; Huang, J.; Pu, J.; Mei, Y.; Nam, K.; Shao, Y. Machine-Learning-Assisted Free Energy Simulation of Solution-Phase and Enzyme Reactions. *J. Chem. Theory Comput.* **2021**, *17*, 5745–5758.
- (19) Gómez-Flores, C. L.; Maag, D.; Kansari, M.; Vuong, V. Q.; Irle, S.; Gräter, F.; Kubař, T.; Elstner, M. Accurate Free Energies for Complex Condensed-Phase Reactions Using an Artificial Neural Network Corrected DFTB/MM Methodology. *J. Chem. Theory Comput.* **2022**, *18*, 1213–1226.

- (20) Komp, E.; Janulaitis, N.; Valleau, S. Progress towards machine learning reaction rate constants. *Phys. Chem. Chem. Phys.* **2022**, *24*, 2692–2705.
- (21) Wu, J.; Shen, L.; Yang, W. Internal force corrections with machine learning for quantum mechanics/molecular mechanics simulations. *J. Chem. Phys.* **2017**, *147*.
- (22) Shen, L.; Yang, W. Molecular Dynamics Simulations with Quantum Mechanics/Molecular Mechanics and Adaptive Neural Networks. *J. Chem. Theory Comput.* **2018**, *14*, 1442–1455.
- (23) Zeng, J.; Giese, T. J.; Ekesan, Ş.; York, D. M. Development of Range-Corrected Deep Learning Potentials for Fast, Accurate Quantum Mechanical/Molecular Mechanical Simulations of Chemical Reactions in Solution. *J. Chem. Theory Comput.* **2021**, *17*, 6993–7009.
- (24) Jorner, K.; Brinck, T.; Norrby, P. O.; Buttar, D. Machine learning meets mechanistic modelling for accurate prediction of experimental activation energies. *Chem. Sci.* **2021**, *12*, 1163–1175.
- (25) Gastegger, M.; Schütt, K. T.; Müller, K. R. Machine learning of solvent effects on molecular spectra and reactions. *Chem. Sci.* **2021**, *12*, 11473–11483.
- (26) Farrar, E. H.; Grayson, M. N. Machine learning and semi-empirical calculations: a synergistic approach to rapid, accurate, and mechanism-based reaction barrier prediction. *Chem. Sci.* **2022**, *13*, 7594–7603.
- (27) Schwaller, P.; Probst, D.; Vaucher, A. C.; Nair, V. H.; Kreutter, D.; Laino, T.; Raymond, J. L. Mapping the space of chemical reactions using attention-based neural networks. *Nat. Mach. Intell.* **2021**, *3*, 144–152.
- (28) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; Palmer, A.; Settels, V.; Jaakkola, T.; Jensen, K.; Barzilay, R. Analyzing Learned Molecular Representations for Property Prediction. *J. Chem. Inf. Model.* **2019**, *59*, 3370–3388.
- (29) Feinberg, E. N.; Joshi, E.; Pande, V. S.; Cheng, A. C. Improvement in ADMET Prediction with Multitask Deep Featurization. *J. Med. Chem.* **2020**, *63*, 8835–8848.
- (30) Klamt, A. Conductor-like Screening Model for Real Solvents: A New Approach to the Quantitative Calculation of Solvation Phenomena. *J. Phys. Chem.* **1995**, *99*, 2224–2235.
- (31) Chung, Y.; Green, W. H. Computing kinetic solvent effects and liquid phase rate constants using quantum chemistry and COSMO-RS methods. *J. Phys. Chem. A* **2023**, [accepted].
- (32) Chung, Y.; Gillis, R. J.; Green, W. H. Temperature-dependent vapor–liquid equilibria and solvation free energy estimation from minimal data. *AIChE J.* **2020**, *66*, e16976.

- (33) Grambow, C. A.; Pattanaik, L.; Green, W. H. Reactants, products, and transition states of elementary chemical reactions based on quantum chemistry. *Sci. Data* **2020**, *7*.
- (34) Grambow, C. A.; Pattanaik, L.; Green, W. H. Reactants, products, and transition states of elementary chemical reactions based on quantum chemistry (1.0.1) [Data set]. *Zenodo* **2020**, DOI: 10.5281/zenodo.3715478.
- (35) Grambow, C. A. Reactants, products, and transition states of radical reactions (1.0.0) [Data set]. *Zenodo* **2020**, DOI: 10.5281/zenodo.3731554.
- (36) Harms, N. D.; Underkoffler, C. E.; West, R. H. Advances in automated transition state theory calculations: improvements on the AutoTST framework. *ChemRxiv* **2020**, DOI: 10.26434/chemrxiv.13277870.v2.
- (37) Perdew, J. P. Density-functional approximation for the correlation energy of the inhomogeneous electron gas. *Phys. Rev. B* **1986**, *33*, 8822–8824.
- (38) Becke, A. D. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **1988**, *38*, 3098–3100.
- (39) Rappoport, D.; Furche, F. Property-optimized Gaussian basis sets for molecular response calculations. *J. Chem. Phys.* **2010**, *133*, 134105.
- (40) Klamt, A.; Diedenhofen, M. A refined cavity construction algorithm for the conductor-like screening model. *J. Comput. Chem.* **2018**, *39*, 1648–1655.
- (41) Klamt, A.; Jonas, V.; Bürger, T.; Lohrenz, J. C. W. Refinement and Parametrization of COSMO-RS. *J. Phys. Chem. A* **1998**, *102*, 5074–5085.
- (42) Eckert, F.; Klamt, A. Fast solvent screening via quantum chemistry: COSMO-RS approach. *AIChE J.* **2002**, *48*, 369–385.
- (43) *TURBOMOLE*, version 7.5; a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH; TURBOMOLE GmbH: Karlsruhe, Germany, 2020. <https://www.turbomole.org> (accessed May 26, 2023).
- (44) Balasubramani, S. G.; Chen, G. P.; Coriani, S.; Diedenhofen, M.; Frank, M. S.; Franzke, Y. J.; Furche, F.; Grotjahn, R.; Harding, M. E.; Hättig, C.; Hellweg, A.; Helmich-Paris, B.; Holzer, C.; Huniar, U.; Kaupp, M.; Marefat Khah, A.; Karbalaeei Khani, S.; Müller, T.; Mack, F.; Nguyen, B. D.; Parker, S. M.; Perlt, E.; Rappoport, D.; Reiter, K.; Roy, S.; Rückert, M.; Schmitz, G.; Sierka, M.; Tapavicza, E.; Tew, D. P.; van Wüllen, C.; Voora, V. K.; Weigend, F.; Wodyński, A.; Yu, J. M. TURBOMOLE: Modular program suite for ab initio quantum-chemical and condensed-matter simulations. *J. Chem. Phys.* **2020**, *152*, 184107.
- (45) Dassault Systèmes, BIOVIA COSMOtherm, Release 2020. <http://www.3ds.com>, (accessed May 26, 2023).

- (46) Dassault Systèmes, BIOVIA COSMObase, Release 2020. <http://www.3ds.com>, (accessed May 26, 2023).
- (47) Lin, Y. S.; Li, G. D.; Mao, S. P.; Chai, J. D. Long-range corrected hybrid density functionals with improved dispersion corrections. *J. Chem. Theory Comput.* **2013**, *9*, 263–272.
- (48) Weigend, F. Accurate Coulomb-fitting basis sets for H to Rn. *Phys. Chem. Chem. Phys.* **2006**, *8*, 1057–1065.
- (49) Zhao, Y.; Truhlar, D. G. The M06 suite of density functionals for main group thermochemistry, thermochemical kinetics, noncovalent interactions, excited states, and transition elements: Two new functionals and systematic testing of four M06-class functionals and 12 other functionals. *Theor. Chem. Acc.* **2008**, *120*, 215–241.
- (50) Dunning, T. H. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* **1989**, *90*, 1007–1023.
- (51) O’Boyle, N. M.; Banck, M.; James, C. A.; Morley, C.; Vandermeersch, T.; Hutchison, G. R. Open Babel: An open chemical toolbox. *J. Cheminform.* **2011**, *3*, 33.
- (52) Spiekermann, K.; Pattanaik, L.; Green, W. H. High accuracy barrier heights, enthalpies, and rate coefficients for chemical reactions. *Sci. Data* **2022**, *9*.
- (53) Varnek, A.; Fourches, D.; Hoonakker, F.; Solov’ev, V. P. Substructural fragments: an universal language to encode reactions, molecular and supramolecular structures. *J. Comput. Aided Mol. Des.* **2005**, *19*, 693–703.
- (54) Hoonakker, F.; Lachiche, N.; Varnek, A.; Wagner, A. A representation to apply usual data mining techniques to chemical reactions - Illustration on the rate constant of S_N2 reactions in water. *Int. J. Artif. Intell. Tools* **2011**, *20*, 253–270.
- (55) Landrum, G. RDKit: Open-Source Cheminformatics. 2006; <http://www.rdkit.org/>, (accessed June 16, 2023).
- (56) Bergstra, J.; Yamins, D.; Cox, D. D. Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. Proceedings of the 30th International Conference on Machine Learning. Atlanta, Georgia, 2013; pp I-115 to I-23.
- (57) Chung, Y.; Vermeire, F. H.; Wu, H.; Walker, P. J.; Abraham, M. H.; Green, W. H. Group Contribution and Machine Learning Approaches to Predict Abraham Solute Parameters, Solvation Free Energy, and Solvation Enthalpy. *J. Chem. Inf. Model.* **2022**, *62*, 433–446.
- (58) Vermeire, F. H.; Green, W. H. Transfer learning for solvation free energies: From quantum chemistry to experiments. *Chem. Eng. J.* **2021**, *418*, 129307.

- (59) Guan, Y.; Coley, C. W.; Wu, H.; Ranasinghe, D.; Heid, E.; Struble, T. J.; Pattanaik, L.; Green, W. H.; Jensen, K. F. Regio-selectivity prediction with a machine-learned reaction representation and on-the-fly quantum mechanical descriptors. *Chem. Sci.* **2021**, *12*, 2198–2208.
- (60) Abraham, M. H.; Acree, W. E. Correlation and prediction of partition coefficients between the gas phase and water, and the solvents dodecane and undecane. *New J. Chem.* **2004**, *28*, 1538–1543.
- (61) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; Duchesnay, E. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
- (62) Stuyver, T.; Coley, C. W. Quantum chemistry-augmented neural networks for reactivity prediction: Performance, generalizability, and explainability. *J. Chem. Phys.* **2022**, *156*.
- (63) Pathak, Y.; Mehta, S.; Priyakumar, U. D. Learning Atomic Interactions through Solvation Free Energy Prediction Using Graph Neural Networks. *J. Chem. Inf. Model.* **2021**, *61*, 689–698.
- (64) Lim, H.; Jung, Y. J. MLSolvA: solvation free energy prediction from pairwise atomistic interactions by machine learning. *J. Cheminform.* **2021**, *13*, 1–10.
- (65) Zhang, D.; Xia, S.; Zhang, Y. Accurate Prediction of Aqueous Free Solvation Energies Using 3D Atomic Feature-Based Graph Neural Network with Transfer Learning. *J. Chem. Inf. Model.* **2022**, *62*, 1840–1848.
- (66) Vermeire, F. H.; Chung, Y.; Green, W. H. Predicting Solubility Limits of Organic Solutes for a Wide Range of Solvents and Temperatures. *J. Am. Chem. Soc.* **2022**, *144*, 10785–10797.
- (67) Litwinienko, G.; Beckwith, A. L.; Ingold, K. U. The frequently overlooked importance of solvent in free radical syntheses. *Chem. Soc. Rev.* **2011**, *40*, 2157–2163.
- (68) Reichardt, C.; Welton, T. *Solvents and Solvent Effects in Organic Chemistry*, fourth edition ed.; Wiley-VCH Verlag GmbH & Co. KGaA: Weinheim, 2011.
- (69) Reuther, A.; Kepner, J.; Byun, C.; Samsi, S.; Arcand, W.; Bestor, D.; Bergeron, B.; Gadepally, V.; Houle, M.; Hubbell, M.; Jones, M.; Klein, A.; Milechin, L.; Mullen, J.; Prout, A.; Rosa, A.; Yee, C.; Michaleas, P. Interactive supercomputing on 40,000 cores for machine learning and data analysis. 2018 IEEE High Performance extreme Computing Conference (HPEC). 2018; pp 1–6.

Graphical TOC Entry

