

Regression-clustering for Improved Accuracy and Training Cost with Molecular-Orbital-Based Machine Learning

Lixue Cheng,[†] Nikola B. Kovachki,[‡] Matthew Welborn,[†] and Thomas F. Miller

III*,[†]

[†] *Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125, USA*

[‡] *Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA*

E-mail: tfm@caltech.edu

Abstract

Machine learning (ML) in the representation of molecular-orbital-based (MOB) features has been shown to be an accurate and transferable approach to the prediction of post-Hartree-Fock correlation energies. Previous applications of MOB-ML employed Gaussian Process Regression (GPR), which provides good prediction accuracy with small training sets; however, the cost of GPR training scales cubically with the amount of data and becomes a computational bottleneck for large training sets. In the current work, we address this problem by introducing a clustering/regression/classification implementation of MOB-ML. In a first step, regression clustering (RC) is used to partition the training data to best fit an ensemble of linear regression (LR) models; in a second step, each cluster is regressed independently, using either LR or GPR; and in a third step, a random forest classifier (RFC) is trained for the prediction of cluster assignments based on MOB feature values. Upon inspection, RC is

found to recapitulate chemically intuitive groupings of the frontier molecular orbitals, and the combined RC/LR/RFC and RC/GPR/RFC implementations of MOB-ML are found to provide good prediction accuracy with greatly reduced wall-clock training times. For a dataset of thermalized (350 K) geometries of 7211 organic molecules of up to seven heavy atoms (QM7b-T), both RC/LR/RFC and RC/GPR/RFC reach chemical accuracy (1 kcal/mol prediction error) with only 300 training molecules, while providing 35000-fold and 4500-fold reductions in the wall-clock training time, respectively, compared to MOB-ML without clustering. The resulting models are also demonstrated to retain transferability for the prediction of large-molecule energies with only small-molecule training data. Finally, it is shown that capping the number of training datapoints per cluster leads to further improvements in prediction accuracy with negligible increases in wall-clock training time.

1 Introduction

Machine-learning (ML) continues to emerge as a versatile strategy in the chemical sciences, with applications to drug discovery,^{1–5} materials design,^{5–9} and reaction prediction.^{5,10–14} An increasing number of ML methods have focused on the prediction of molecular properties, including quantum mechanical electronic energies,^{15–31} densities,^{26,32–36} and spectra.^{37–41} Most of this work has focused on ML in the representation of atom- or geometry-specific features, although more abstract representations are gaining increased attention.^{42–48}

We recently introduced a rigorous factorization of the post-Hartree-Fock correlation energy into contributions from pairs of occupied molecular orbitals and showed that these pair contributions could be compactly represented in the space of molecular-orbital-based (MOB) features to allow for straightforward ML regression.^{47,48} This MOB-ML method was demonstrated to accurately predict second-order Møller-Plessett perturbation theory (MP2)^{49,50} and coupled cluster with singles, doubles and perturbative triples (CCSD(T))^{51,52} energies of different benchmark systems, including the QM7b-T and GDB-13-T datasets of thermalized drug-like organic molecules. While providing good accuracy with a modest amount of training data, the accuracy of MOB-ML

in these initial studies was limited by the high computational cost ($\mathcal{O}(N^3)$) of applying Gaussian Process Regression (GPR) to the full set of training data.⁴⁸

In this work, we combine MOB-ML with regression clustering (RC) to overcome this bottleneck in computational cost and accuracy. The training data are clustered via RC to discover locally linear structures. By independently regressing these subsets of the data, we obtain MOB-ML models with greatly reduced training costs while preserving prediction accuracy and transferability.

2 Theory

2.1 Molecular-orbital based machine learning (MOB-ML)

The MOB-ML method is based on the observation that the correlation energy for any post-Hartree-Fock wavefunction theory can be exactly decomposed as a sum over occupied molecular orbitals (MOs) via Nesbet’s theorem,^{53,54}

$$E_c = \sum_{ij}^{\text{occ}} \epsilon_{ij}, \quad (1)$$

where E_c is the correlation energy and ϵ_{ij} is the pair correlation energy corresponding to occupied MOs i and j . The pair correlation energies can be expressed as a functional of the set of (occupied and unoccupied) MOs, appropriately indexed by i and j , such that

$$\epsilon_{ij} = \epsilon \left[\{ \phi_p \}^{ij} \right]. \quad (2)$$

The functional ϵ maps the Hartree-Fock MOs to the pair correlation energy, regardless of the molecular composition or geometry, such that it is a universal functional for all chemical systems. To bypass the expensive post-Hartree-Fock evaluation procedure, MOB-ML approximates ϵ_{ij} by machine learning two functionals, ϵ_d^{ML} and ϵ_o^{ML} , which correspond to diagonal and off-diagonal terms of the sum in Eq. 1.

$$\epsilon_{ij} \approx \begin{cases} \epsilon_{\text{d}}^{\text{ML}}[\mathbf{f}_i] & \text{if } i = j \\ \epsilon_{\text{o}}^{\text{ML}}[\mathbf{f}_{ij}] & \text{if } i \neq j \end{cases} \quad (3)$$

The MOB-ML feature vectors \mathbf{f}_i and \mathbf{f}_{ij} are comprised of unique elements of the Fock, Coulomb and exchange matrices between ϕ_i , ϕ_j , and the set of virtual orbitals. Without loss of generality, we perform MOB-ML using localized MOs (LMOs) to improve transferability across chemical systems.⁴⁷ Detailed descriptions of feature design are provided in our previous work,^{47,48} and the features employed here are unchanged from those detailed in Ref. 48.

2.2 Local linearity of MOB feature space

It has been previously emphasized that MOB-ML facilitates transferability across chemical systems, even allowing for predictions involving molecules with elements that do not appear in the training set,⁴⁷ due to the fact that MOB features provide a compact and highly abstracted representation of the electronic structure. However, it is worth additionally emphasizing that this transferability benefits from the smooth variation and local linearity of the pair correlation energies as a function of MOB feature values associated with different molecular geometries and even different molecules.

Figure 1 illustrates these latter properties for a σ -bonding orbital in a series of simple molecules. On the y-axis, we plot the diagonal contribution to the correlation energy associated with this orbital (ϵ_{ii}), computed at the MP2/cc-pvTZ level of theory. On the x-axis, we plot the value of a particular MOB feature, the Fock matrix element for the that localized orbital, F_{ii} . For each molecule, a range of geometries is sampled from the Boltzmann distribution at 350 K, with each plotted point corresponding to a different sampled geometry.

It is immediately clear from the figure that the pair correlation energy varies smoothly and linearly as a function of the MOB feature value. Moreover, the slope of the linear curve is remarkably consistent across molecules. This illustration suggests that MOB features may lead to accurate regression of correlation energies using simple machine learning models (even linear models), and it

also indicates the basis for the robust transferability of MOB-ML across diverse chemical systems, including those with elements that do not appear in the training set.

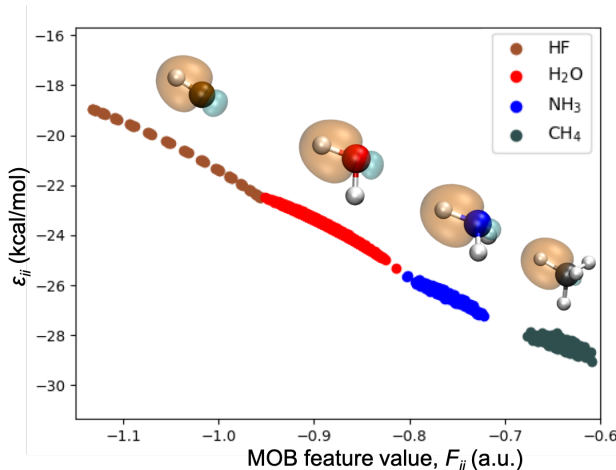


Figure 1: The diagonal pair correlation energy (ϵ_{ii}) for a localized σ -bond in four different molecules at thermally sampled geometries (at 350 K), computed at the MP2/cc-pvTZ level of theory. The diagonal pair correlation energies for HF, NH₃, and CH₄ are shifted vertically downward relative to those of HF by 3.407, 6.289, and 7.772 kcal/mol for H₂O, NH₃ and CH₄. Illustrative σ -bond LMOs are shown for each molecule.

2.3 Regression clustering with a greedy algorithm

To take advantage of the local linearity of pair correlation energies as a function of MOB features, we propose a strategy to discover optimally linear clusters using regression clustering (RC).⁵⁵ Consider the set of M datapoints $\{\mathbf{f}_t, \epsilon_t\} \subset \mathbb{R}^d \times \mathbb{R}$, where d is the length of the MOB feature vector and where each datapoint is indexed by t and corresponds to a MOB feature vector and the associated reference value (i.e., label) for the pair correlation energy. To separate these datapoints into locally linear clusters, S_1, \dots, S_N , we seek a solution to the optimization problem

$$\min_{S_1, \dots, S_N} \sum_{k=1}^N \sum_{t \in S_k} |A(S_k) \cdot \mathbf{f}_t + b(S_k) - \epsilon_t|^2 \quad (4)$$

where $A(S_k) \in \mathbb{R}^d$ and $b(S_k) \in \mathbb{R}$ are obtained via ordinary least squares (OLS) solution,

$$\begin{bmatrix} \mathbf{f}_{t_1}^T & 1 \\ \vdots & \vdots \\ \mathbf{f}_{t_{|S_k|}}^T & 1 \end{bmatrix} \begin{bmatrix} A(S_k) \\ b(S_k) \end{bmatrix} = \begin{bmatrix} \varepsilon_{t_1} \\ \vdots \\ \varepsilon_{t_{|S_k|}} \end{bmatrix}. \quad (5)$$

Each resulting S_k is the set of indices t assigned to cluster k comprised of $|S_k|$ datapoints. To perform the optimization in Eq. 4, we employ a modified version of the greedy algorithm proposed in Ref. 56 (Algorithm 1). In general, solutions to Eq. 4 may overlap, such that $S_k \cap S_l \neq \emptyset$ for $k \neq l$; however, the proposed algorithm enforces that clusters remain pairwise-disjoint.

Algorithm 1 Greedy algorithm for the solution of Eq. 4.

Input: Initial clusters: S_1, \dots, S_N

Output: Data clusters S_1, \dots, S_N

```

1: for  $k \leftarrow 1$  to  $N$  do
2:    $A(S_k), b(S_k) \leftarrow$  OLS solution of Eq. 5
3: end for
4: while not converged do
5:   for  $k \leftarrow 1$  to  $N$  do
6:      $S_k \leftarrow \{t \in \{1, \dots, M\} : \arg \min_{n \in \{1, \dots, N\}} |A(S_n) \cdot \mathbf{f}_t + b(S_n) - \varepsilon_t|^2 = k\}$ 
7:   end for
8:   for  $k \leftarrow 1$  to  $N$  do
9:      $A(S_k), b(S_k) \leftarrow$  OLS solution of Eq. 5
10:  end for
11: end while
```

Algorithm 1 has a per-iteration runtime of $\mathcal{O}(Md^2)$, since we compute N OLS solutions each with runtime $\mathcal{O}(|S_k|d^2)$ and since $\sum_{k=1}^N |S_k| = M$. However, the algorithm can be trivially parallelized to reach a runtime of $\mathcal{O}(\max(|S_k|)d^2)$. A key operational step in this algorithm is line 6, which can be explained in simple terms as follows: we assign each datapoint, indexed by t , to the cluster to which it is closest, as measured by the squared linear regression distance metric,

$$|D_{n,t}|^2 = |A(S_n) \cdot \mathbf{f}_t + b(S_n) - \varepsilon_t|^2 \quad (6)$$

where $D_{n,t}$ is the distance of this point to cluster n . In principle, a datapoint could be equidistant to two or more different clusters by this metric; in such cases, we randomly assign the datapoint to only one of those equidistant clusters to enforce the pairwise-disjointness of the resulting clusters. Convergence of the greedy algorithm is measured by the decrease in the objective function of Eq. 4.

Figure 2 illustrates RC in a simple one-dimensional example for which unsupervised clustering approaches will fail to reveal the underlying linear structure. To create two clusters of nearly linear data that overlap in feature space, the interval of feature values on $[0, 1]$ is uniformly discretized, such that $\mathbf{f}_t = (t - 1)/(M - 1)$ for $t = 1, \dots, M$. Then, $M/2$ of the feature values are randomly chosen without replacement for cluster S_1 while the remainder are placed in S_2 ; the energy labels associated with each feature value are then generated using

$$\varepsilon_t = \mathbf{f}_t + \xi_{t,1}, \quad t \in S_1$$

and

$$\varepsilon_t = -\mathbf{f}_t + 1 + \xi_{t,2}, \quad t \in S_2$$

where $\xi_{t,k} \sim \mathcal{N}(0, 0.1^2)$ is an i.i.d. sequence. The resulting dataset is shown in Fig. 2a.

Application of the RC method to this example is illustrated in Figs. 2(b-d). The greedy algorithm is initialized by randomly assigning each datapoint to either S_1 or S_2 (Fig. 2b). Then, with only a small number of iterations (Figs. 2c and d), the algorithm converges to clusters that reflect the underlying linear character. For comparison, Fig. 2e shows the clustering that is obtained upon convergence of the standard K-means algorithm,⁵⁷ initialized with random cluster assignments. Unlike RC, the K-means algorithm prioritizes the compactness of clusters, resulting in a final clustering that is far less amenable to simple regression. While we recognize that the correct clustering could potentially be obtained using K-means when the dimensions of \mathbf{f}_t and ε_t are comparable, this is not the case for MOB-ML applications since \mathbf{f}_t is typically at least 10-dimensional and ε_t is a scalar; the RC approach does not suffer from this issue. Finally, we have confirmed that initialization of RC from the clustering in Fig. 2e rapidly returns to the results in Fig. 2d, requiring only a

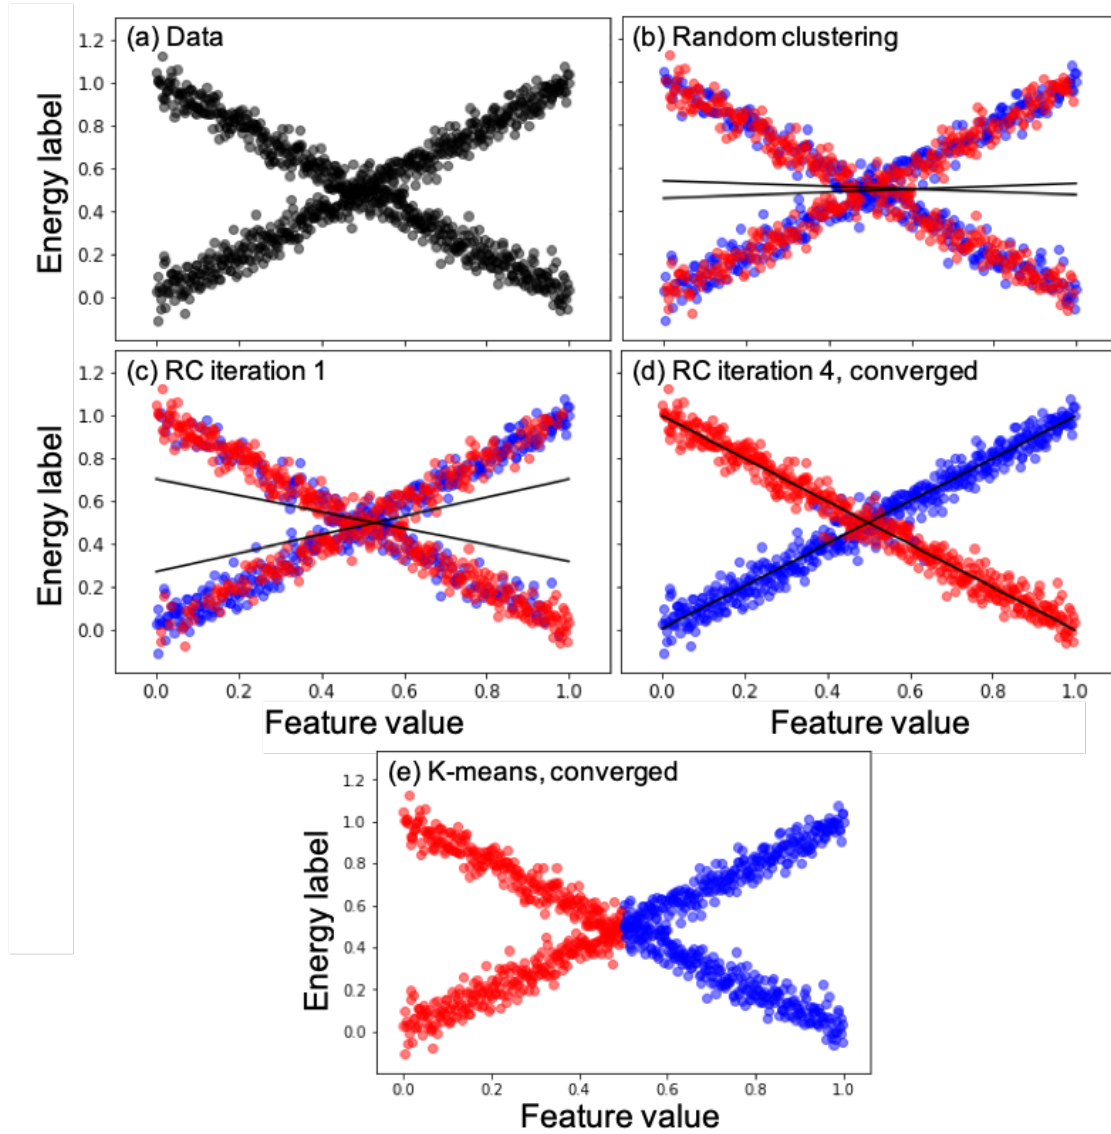


Figure 2: Comparison of clustering algorithms for (a) a dataset comprised of two cluster of nearly linear data that overlap in feature space, using (b-d) RC and (e) standard K-means clustering. (b) Random initialization of the clusters for the greedy algorithm, with datapoint color indicating cluster assignment. (c) Cluster assignments after one iteration of the greedy algorithm. (d) Converged cluster assignments after four iterations of the greedy algorithm. For panels (b-d), two linear regression lines at each iteration are shown in black. (e) Converged cluster assignments obtained using K-means clustering, which fails to reveal the underlying linear structure of the clusters.

couple of iterations of the greedy algorithm.

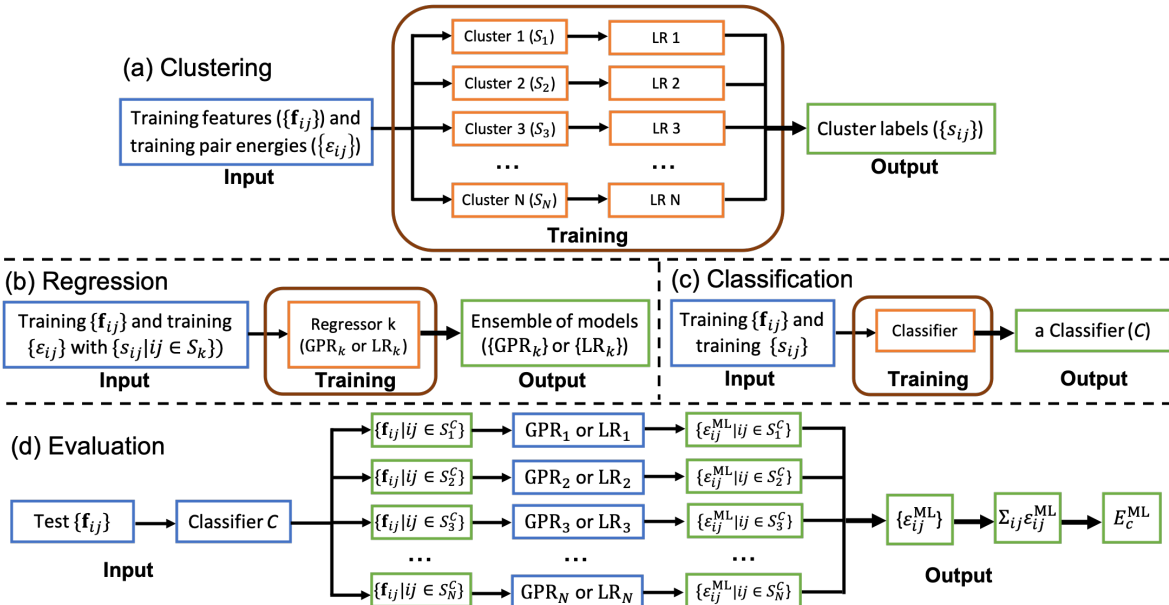


Figure 3: The MOB-ML clustering/regression/classification workflow. (a) Clustering of the training dataset of MOB-ML feature vectors and energy labels using RC to obtain optimized linear clusters and to provide the cluster labels for the feature vectors. (b) Regression of each cluster of training data (using LR or GPR), to obtain the ensemble of cluster-specific regression models. (c) Training a classifier (RFC) from the MOB-ML feature vectors and cluster labels for the training data. (d) Evaluating the predicted MOB-ML pair correlation energy from a test feature vector is performed by first classifying the feature vector into one of the clusters, then evaluating the cluster-specific regression model. In each panel, blue boxes indicate input quantities, orange boxes indicate training intermediates, and green boxes indicate the resulting labels, models and pair correlation energy predictions.

3 Calculation Details

Results are presented for QM7b-T,⁴⁸ a thermalized version of the QM7b set⁵⁸ of 7211 molecules with up to seven C, O, N, S, and Cl heavy atoms, as well as for GDB-13-T,⁴⁸ a thermalized version of the GDB-13 set⁵⁹ of molecules with thirteen C, O, N, S, and Cl heavy atoms. The MOB-ML features employed in the current study are identical to those previously provided.⁴⁸ Reference pair correlation energies are computed using MP2⁴⁹ and using CCSD(T).^{51,52} The MP2 reference data were obtained with the cc-pVTZ basis set,⁶⁰ whereas the CCSD(T) data were obtained using the

cc-pVDZ basis set.⁶⁰ All employed training and test datasets are provided in Ref. 48.

3.1 Regression Clustering (RC)

RC is performed using the ordinary least square linear regression implementation in the SCIKIT-LEARN package.⁶¹ Unless otherwise specified, we initialize the greedy algorithm from the results of K-means clustering, also implemented in SCIKIT-LEARN; K-means initialization was found to improve the subsequent training of the random forest classifier (RFC) in comparison to random initialization. It is found that neither L1 nor L2 regularization had significant effect on the rate of convergence of the greedy algorithm, so neither is employed in the results presented here. It is found that a convergence threshold of 1×10^{-8} kcal²/mol² for the loss function of the greedy algorithm (Eq. 4) leads to no degradation in the final MOB-ML regression accuracy (Fig. S2); this value is employed throughout.

3.2 Regression

Two different regression models are employed in the current work. The first is ordinary least-squares linear regression (LR), as implemented in SCIKIT-LEARN. The second is Gaussian Process Regression, as implemented in the GPY 1.9.6 software package.⁶² Regression is independently performed for the training data associated with each cluster, yielding a local regression model for each cluster. Also, as in our previous work,^{47,48} regression is independently performed for the diagonal and off-diagonal pair correlation energies (ϵ_d^{ML} and ϵ_o^{ML}) yielding independent regression models for each (Eq. 3).

GPR is performed using a negative log marginal likelihood objective. As in our previous work,⁴⁸ the Matérn 5/2 kernel is used for regression of the diagonal pair correlation energies and the Matérn 3/2 kernel is used for the off-diagonal pair correlation energies; in both cases, white noise regularization⁶³ is employed, and the GPR is initialized with unit lengthscale and variance.

3.3 Classification

An RFC is trained on MOB-ML features and cluster labels for a training set and then used to predict the cluster assignment of test datapoints in MOB-ML feature space. We employ the RFC implementation in SCIKIT-LEARN, using with 200 trees, the entropy split criteria,⁶⁴ and balanced class weights.⁶⁴ Alternative classifiers were also tested in this work, including K-means, Linear SVM,⁶⁵ and AdaBoost;⁶⁶ however, these schemes were generally found to yield less accurate MOB-ML energy predictions than RFC.

For comparison, a “perfect” classifier is obtained by simply including the test data within the RC training set. While useful for the analysis of prediction errors due to classification, this scheme is not generally practical because it assumes prior knowledge of the reference energy labels for the test molecules. Since the perfect classifier avoids mis-classification of the test data by construction, it should be regarded as a best case scenario for the performance of the clustering/regression/classification approach.

3.4 The clustering/regression/classification workflow

Fig. 3 summarizes the combined work flow for training and evaluating a MOB-ML model with clustering. The training involves three steps: First, the training dataset of MOB-ML feature vectors and energy labels are assigned to clusters using the RC method (panel a). Second, for each cluster of training data, the regression model (LR or GPR) is trained, to enable the prediction of pair correlation energies from the MOB-ML vector. Third, a classifier is trained from the MOB-ML feature vectors and cluster labels for the training data, to enable the prediction of the cluster assignment from a MOB-ML feature vector.

The resulting MOB-ML model is specified in terms of the method of clustering (RC, for all results presented here), the method of regression (either LR or GPR), and the method of classification (either RFC or the perfect classifier). In referring to a given MOB-ML model, we employ a notation that specifies these options (e.g., RC/LR/RFC or RC/GPR/perfect).

Evaluation of the trained MOB-ML model is explained in Fig. 3d. A given molecule is first

decomposed into a set of test feature vectors associated with the pairs of occupied MOs. The classifier is then used to assign each feature vector to an associated cluster. The cluster-specific regression model is then used to predict the pair correlation energy from the MOB feature vector. And finally, the pair correlation energies are summed to yield the total correlation energy for the molecule.

To improve the accuracy and reduce the uncertainty in the MOB-ML predictions, ensembles of 10 independent models using the clustering/regression/classification workflow are trained, and the predictive mean and the corresponding standard error of the mean (SEM) are computed by averaging over the 10 models; a comparison between the learning curves⁶⁷ from a single run and from averaging over the 10 independent models is included in Supporting Information Fig. S1. As described here, the predicted correlation energies may exhibit discontinuities as a function of nuclear position, due to changes in the assignment of feature vectors among the clusters; moving forward, this may be avoided with the use of soft (or fuzzy) clustering algorithms.⁶⁸

4 Results

4.1 Clustering and classification in MOB feature space

We begin by showing that the situation explored in Fig. 2, in which locally linear clusters overlap, also arises in realistic chemical applications of MOB-ML. We consider the QM7b-T set of drug-like molecules with thermalized geometries, using the diagonal pair correlation energies ϵ_d^{ML} computed at the MP2/cc-pVTZ level. Randomly selecting 1000 molecules for training, we perform RC on the dataset comprised of these energy labels and feature vectors, using $N = 20$ optimized clusters; the sensitivity of RC to the choice of N is examined later.

In many cases, the resulting clusters are well separated, such that the datapoints for one cluster have small distances (as measured by the linear regression distance metric, Eq. 6) to the cluster which it belongs to and large distances to all other clusters. However, the clusters can also overlap. Fig. 4a illustrates this overlap for two particular clusters (labeled 1 and 2) obtained from the QM7b-

T diagonal-pair training data.

Each datapoint assigned to cluster 1 (blue) is plotted according to its distance to both cluster 1 and cluster 2; likewise for the datapoints in cluster 2 (red). The datapoints for which the distances to both clusters approach zero correspond to regions of overlap between the clusters in the high-dimensional space of MOB-ML features, akin to the case shown in Fig. 2.

Finally, in Fig. 4b, we illustrate the classification of the feature vectors into clusters. An RFC is trained on the feature vectors and cluster labels for the diagonal pairs of 1000 QM7b-T molecules in the training set, and the classifier is then used to predict the cluster assignment for the feature vectors associated with the remaining diagonal pairs of 6211 molecules in QM7b-T. For clusters 1 and 2, we then analyze the accuracy of the RFC by plotting the linear regression distance for each datapoint to the two clusters, as well as indicating the RFC classification of the feature vector. Each red datapoint in Fig. 4b that lies above the diagonal line of reflection is mis-classified into cluster 2, and similarly, each blue datapoint that lies below the line of reflection is mis-classified into cluster 1. The figure illustrates that while RFC is not a perfect means of classification, it is at least qualitatively correct. Later, in the results section, we will analyze the sources of MOB-ML prediction errors due to mis-classification by comparing energy predictions obtained with perfect classification versus RFC.

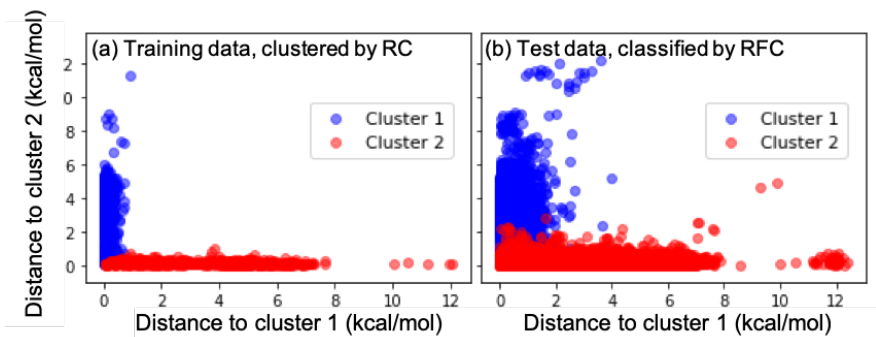


Figure 4: (a) Illustration of the overlap of clusters obtained via RC for the training set molecules from QM7b-T. (b) Classification of the datapoints for the remaining test molecules from QM7b-T, using RFC. Distances correspond to the linear regression metric defined in Eq. 6.

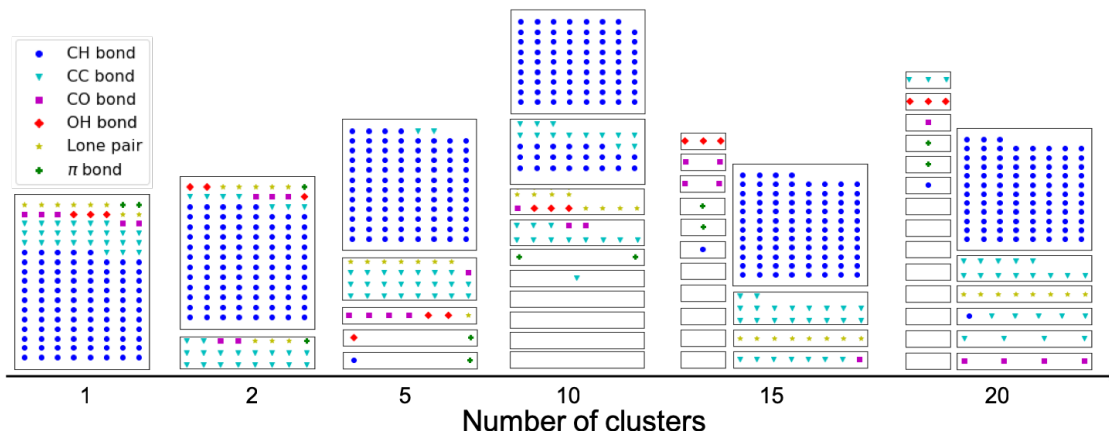


Figure 5: Analyzing the results of clustering/classification in terms of chemical intuition. Using a training set of 500 randomly selected molecules from QM7b-T, RC is performed for the diagonal pair correlation energies, ϵ_d^{ML} , with a range of cluster numbers, N , and for each clustering, an RFC is trained. Then, the trained classifier is applied to a set of test molecules (CH_4 , C_2H_6 , C_2H_4 , C_3H_8 , $\text{CH}_3\text{CH}_2\text{OH}$, CH_3OCH_3 , $\text{CH}_3\text{CH}_2\text{CH}_2\text{CH}_3$, $\text{CH}_3\text{CH}(\text{CH}_3)\text{CH}_3$, $\text{CH}_3\text{CH}_2\text{CH}_2\text{CH}_2\text{CH}_2\text{CH}_2\text{CH}_3$, $(\text{CH}_3)_3\text{CCH}_2\text{OH}$, and $\text{CH}_3\text{CH}_2\text{CH}_2\text{CH}_2\text{CH}_2\text{CH}_2\text{OH}$) which have chemically intuitive LMO types, as indicated in the legend. The LMOs are successfully resolved according to type by the classifier as N increases. Empty boxes correspond to clusters into which none of the LMOs from the test set is classified; these are expected since the training set is more diverse than the test set.

4.2 Chemically intuitive clusters

To address this, we employ a training set of 500 randomly selected molecules from QM7b-T, and we perform regression clustering for the diagonal pair correlation energies ϵ_d^{ML} with a range of total cluster numbers, up to $N = 20$. For each clustering, we then train an RFC. Finally, each trained RFC is independently applied to a set of test molecules with easily characterized valence molecular orbitals (listed in the caption of Fig. 5), to see how the feature vectors associated with valence occupied LMOs are classified among the optimized clusters.

Figure 5 presents the results of this exercise, clearly indicating the agreement between chemical intuition and the predictions of the RFC. As the number of clusters increases, the feature vectors associated with different valence LMO types are resolved into different clusters; and with a sufficiently large number of clusters (15 or 20), each cluster is dominated by a single type of LMO while each LMO type is assigned to a small number of different clusters. The empty boxes in Fig. 5 reflect that the training set contains a larger diversity of LMO types than the 11 test

molecules, which is expected. The observed consistency of the clustering/classification method presented here with chemical intuition is of course promising for the accurate local regression of pair correlation energies, which is the focus of the current work; however, the results of Fig. 5 also suggest that the clustering/classification of chemical systems in MOB-ML feature space provides a powerful and highly general way of mapping the structure of chemical space for other applications, including explorative or active ML applications.⁶⁹

4.3 Sensitivity to the number of clusters

We now explore the sensitivity of the MOB-ML clustering/regression/classification implementation to the number of employed clusters. In particular, we investigate the mean absolute error (MAE) of the MOB-ML predictions for the diagonal ($\sum_i \epsilon_{ii}$) and off-diagonal ($\sum_{i \neq j} \epsilon_{ij}$) contributions to the total correlation energy, as a function of the number of clusters, N , used in the RC. The MOB-ML models employ linear regression and RFC classification (i.e., the RC/LR/RFC protocol); the training set is comprised of 1000 randomly chosen molecules from QM7b-T, and the test set contains the remaining molecules in QM7b-T.

Figure 6 presents the result of this calibration study, plotting the prediction MAE as a function of the number of clusters. Not surprisingly, the prediction accuracy for both the diagonal and off-diagonal contributions improves with N , although it eventually plateaus in both cases. For the diagonal contributions, the accuracy improves most rapidly up to approximately 20 clusters, in accord with the observations in Fig. 5; and for the off-diagonal contributions, a larger number of clusters is useful for reducing the MAE error, which is sensible given the greater variety of feature vectors that can be created from pairs of LMOs rather than only individual LMOs. Appealingly, there does not seem to be a strong indication of MAE increases due to “over-clustering”. While recognizing that the optimal number of clusters will, in general, depend somewhat on the application and the regression method (i.e., LR versus GPR), the results in Fig. 6 nonetheless provide useful guidance with regard to the appropriate values of N . Throughout the remainder of the study, we employ a value of $N = 20$ for the MOB-ML prediction of diagonal contributions to the corre-

lation energy and a value of $N = 70$ for the off-diagonal contributions; however, we recognize that these choices could be further optimized.

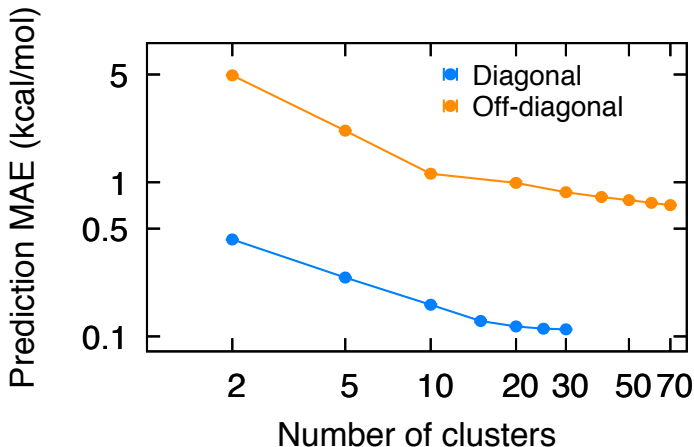


Figure 6: Illustration of the sensitivity of MOB-ML predictions for the diagonal and off-diagonal contributions to the correlation energy for the QM7b-T set of molecules, using a subset of 1000 molecules for training and the RC/LR/RFC protocol. The standard error of the mean (SEM) for the predictions is smaller than the size of the plotted points.

4.4 Performance and training cost of MOB-ML with RC

We now investigate the effect of clustering on the accuracy and training costs of MOB-ML for applications to sets of drug-like molecules. Figure 7a presents learning curves (on a linear-linear scale) for various implementations of MOB-ML applied to MP2/cc-pVTZ correlation energies, with the training and test sets corresponding to non-overlapping subsets of QM7b-T. In addition to the new results obtained using RC, we include the MOB-ML results from our previous work (GPR without clustering).⁴⁸

Figure 7a yields three clear observations. The first is that the use of RC with RFC (i.e., RC/GRP/RFC and RC/LR/RFC) leads to slightly less efficient learning curves than our previous implementation without clustering, at least when efficiency is measured in terms of the number of training molecules. Both the RC/GPR/RFC and RC/LR/RFC protocols require approximately 300 training molecules to reach the 1 kcal/mol per seven heavy atoms threshold for chemical accuracy employed here, whereas MOB-ML without clustering requires approximately half as many train-

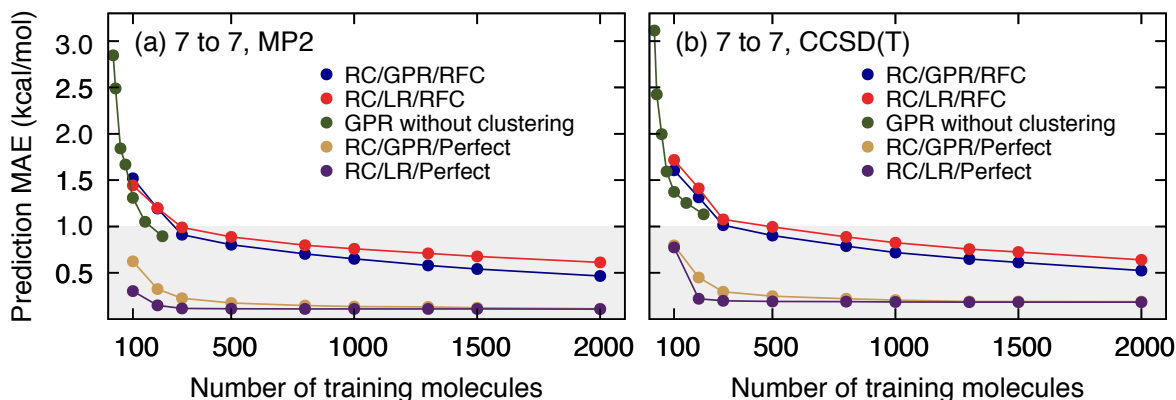


Figure 7: Learning curves for various implementations of MOB-ML applied to (a) MP2/cc-pVTZ and (b) CCSD(T)/cc-pVDZ correlation energies, with the training and test sets corresponding to non-overlapping subsets of the QM7b-T set of drug-like molecules with up to heavy seven atoms. Results obtained using GPR without clustering (green) are reproduced from Ref. 48. The gray shaded area corresponds to a MAE of 1 kcal/mol per seven heavy atoms. The prediction SEM is smaller than the plotted points. The log-log version of this plot is provided in Fig. S3.

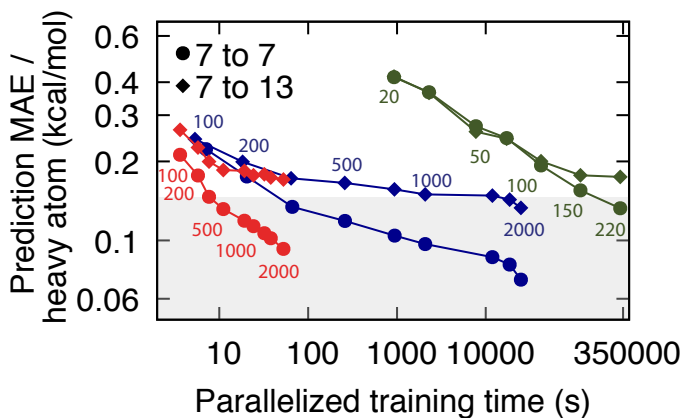


Figure 8: Training costs and transferability of MOB-ML with clustering (RC/LR/RFC, red; RC/GPR/RFC, blue) and without clustering (green, Ref. 48), applied to correlation energies at the MP2/cc-pVTZ level. Prediction errors are plotted as a function of wall-clock training time. Training sets are comprised of subsets of the QM7b-T dataset, with the number of training molecules indicated via datapoint labels. Correlation energy predictions are made for test sets comprised of the remaining seven-heavy-atom molecules from QM7b-T (circles) and the thirteen-heavy-atom molecules from GDB-13-T (diamonds). Both MAE prediction errors and parallelized wall-clock training times are plotted on a log scale. The gray shaded area corresponds to a MAE of 1 kcal/mol per seven heavy atoms. The prediction SEM is smaller than the plotted points. Details of the parallelization and employed computer hardware are described in the text.

ing molecules. The second observation is that the classifier is the dominant source of prediction error in these results. Comparison of results using RFC versus the perfect classifier (which utilizes prior knowledge of the energy labels and this thus not generally practical), reveals a dramatic reduction in the prediction error, regardless of the regression method. This result indicates that there is potentially much to be gained from the development of improved classifiers for MOB-ML applications. A third observation is that with a perfect classifier, the LR slightly outperforms GPR, given that the clusters are optimized to be locally linear; however, GPR slightly outperforms LR in combination with the RFC, indicating that GPR is less sensitive to classification error than LR.

Figure 7b presents the corresponding results at the CCSD(T)/cc-pVDZ level of theory. The same trends emerge as the ones at the MP2/cc-pVTZ level of theory. As seen in previous work, the training efficiency of MOB-ML with respect to the size reference dataset is found to be largely insensitive to the level of electronic structure theory.^{47,48}

Figure 8 explores the training costs and transferability of MOB-ML models that employ RC. In all cases, the models are trained on random subsets of molecules from QM7b-T with up to seven heavy atoms, and predictions are made either on the remaining molecules of QM7b-T (circles) or on the GDB-13-T set (diamonds); it has previously been shown that MOB-ML substantially outperforms the FCHL atom-based-feature method in terms of transferability from small to large molecules.⁴⁸ The parallelization of the training steps are implemented as follows. Within the RC step, the LR for each cluster is performed independently on a different core of a 16-core Intel Skylake (2.1 GHz) CPU processor. Within the regression step, the LR or GPR for each cluster is likewise performed independently on a different core. For RFC training, we apply parallel 200 cores using the parallel implementation of SCIKIT-LEARN, since there are 200 trees. The regression and RFC training are independent of each other and are thus also trivially parallelizable.

Focusing first on the predictions for seven-heavy-atom molecules (circles), it is clear from Fig. 8 that RC leads to large improvements in the efficiency of the MOB-ML wall-clock training costs. Although it requires somewhat more training molecules than MOB-ML without clustering, MOB-ML with clustering enables chemical accuracy to be reached with the training cost reduced

by a factor of approximately 4500 for RC/GPR/RFC and of 35000 for RC/LR/RFC. Remarkably, for predictions within the QM7b-T set, chemical accuracy can be achieved using RC/LR/RFC with a wall-clock training time of only 7.7 s.

Figure 8 also demonstrates the transferability of the MOB-ML models for predictions on the GDB-13-T set of thirteen-heavy-atom molecules (diamonds). In general, it is seen that the degradation in the MAE per atom is greater for the RC/LR/RFC than for RC/GPR/RFC, due to the previously mentioned sensitivity of LR to classification error. However, we note that the RC/GPR/RFC enables predictions on GDB-13-T (blue, diamonds) that meet the per-atom threshold of chemical accuracy used here, whereas that threshold was not achievable without clustering (green, diamonds) due to the prohibitive training costs involved.

The improved efficiency of MOB-ML training with the use of clustering arises from the cubic scaling of standard GPR in terms of training time ($\mathcal{O}(M^3)$, where M is number of training pairs).⁶³ Trivial parallelization over the independent regression of the clusters reduces this training time cost to the cube of largest cluster. We note that other kernel-based ML methods with high complexity in training time, like Kernel Ridge Regression,⁷⁰ would similarly benefit from clustering. For the RC/LR/RFC and RC/GPR/RFC results presented in Fig. 8, a breakdown of the training time contributions for each step of the clustering/regression/classification workflow as a function of the size of the training dataset is shown in Fig. S4; this supporting information figure confirms that the GPR regression dominates the total training (and prediction) costs for the RC/GPR/RFC implementation, whereas training the RFC dominates the training costs for RC/LR/RFC. In addition to improved efficiency in terms of training time, clustering also bring benefits in terms of the memory costs for MOB-ML training, due to the quadratic scaling of GPR memory costs in terms of the size of the dataset.

Finally, returning to the learning curves, we compare the results for MOB-ML both with and without clustering to recent work⁷¹ using Faber-Christensen-Huang-Lilienfeld (FCHL) features. Fig. 9 shows these various learning curves for the MP2/cc-pVTZ correlation energies. For Fig. 9a, the training and test sets correspond to non-overlapping subsets of QM7b-T, and Fig. 9b shows the

transferability of the same models trained using QM7b-T to predict the energies for GDB-13-T. Fig. 9a again shows that MOB-ML RC/GPR/RFC requires slightly more training geometries than MOB-ML without clustering, yet both MOB-ML protocols are more efficient in terms of training data than either the FCHL18⁷² or FCHL19 implementations.⁷¹ Like MOB-ML with clustering, the FCHL19 implementation was developed to reduce training times.

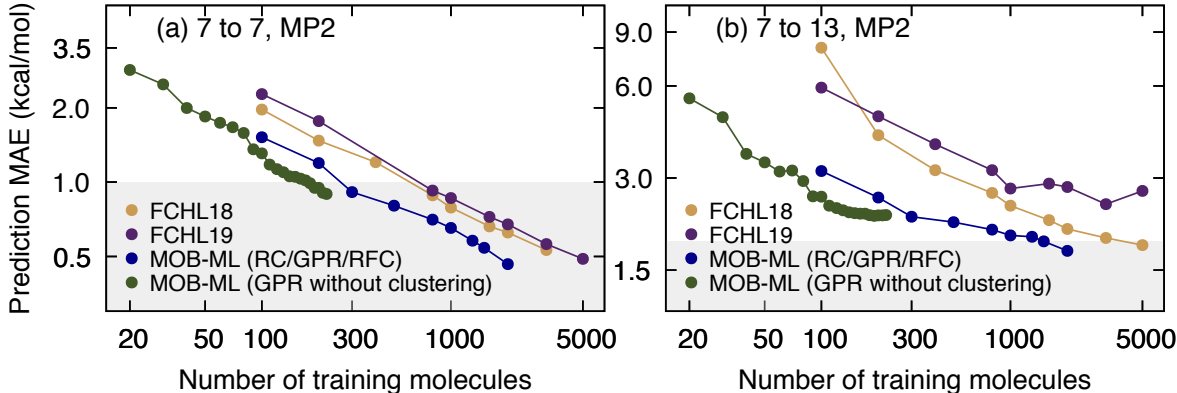


Figure 9: Comparison of learning curves for MP2/cc-pVTZ correlation energies obtained using MOB-ML (with and without clustering) versus FCHL18 and FCHL19. Part (a) presents results for which both the training and test sets include molecules from QM7b-T, and part (b) presents results for which the training set includes molecules from QM7b-T and the test set includes molecules from GDB-13-T. The MAE are plotted on a log-log scale as a function of number of training molecules. The gray shaded area corresponds to a MAE of 1 kcal/mol per seven heavy atoms. Results for FCHL18 and FCHL19 were digitally captured from Ref. 71.

4.5 Capping the cluster size

Since the parallelized training time for RC/GPR/RFC is dominated by the GPR regression of the largest cluster (Fig. S4), a natural question is whether additional computational savings and adequate prediction accuracy could be achieved by simply capping the number of datapoints in the largest cluster. In doing so, we define $S_{\max}^{N_{\text{cap}}}$ to be the number of datapoints in the largest cluster obtained when the RC with the greedy algorithm is applied to a training dataset of N_{cap} molecules from QM7b-T. Upon specifying N_{cap} (and thus $S_{\max}^{N_{\text{cap}}}$), the RC/GPR/RFC implementation is modified as follows. For a given number of training molecules (which will typically exceed N_{cap}), the RC step is performed as normal. However, at the end of the RC step, datapoints for clusters whose size

exceeds $S_{\max}^{N_{\text{cap}}}$ are discarded at random until all clusters contain $S_{\max}^{N_{\text{cap}}}$ or fewer datapoints. The GPR and RFC training steps are performed as before, except using this set of clusters that are capped in size. The precise value of $S_{\max}^{N_{\text{cap}}}$ will vary slightly depending on which training molecules are randomly selected for training and the convergence of the greedy algorithm, but typical values for $S_{\max}^{N_{\text{cap}}}$ are 672, 1218, 1863, 3005, and 4896 for $N_{\text{cap}} = 100, 200, 300, 500$ and 800, respectively, and those values will be used for the numerical tests presented here.

Figure 10a demonstrates that capping the maximum cluster size allows for substantial improvements in accuracy when the number of training molecules exceeds N_{cap} . Specifically, the figure shows the effect of capping on RC/GPR/RFC learning curves for MP2/cc-pVTZ correlation energies, with the training and test sets corresponding to non-overlapping subsets of QM7b-T. As a baseline, note that with 100 training molecules, the RC/GPR/RFC implementation yields a prediction MAE of approximately 1.5 kcal/mol. However, if the maximum cluster size is capped at $N_{\text{cap}} = 100$ and 300 training molecules are employed, then the prediction MAE drops to approximately 1.0 kcal/mol while the parallelized training cost for RC/GPR/RFC will be unchanged so long as it remains dominated by the size of the largest cluster. As expected, Fig. 10a shows that the learning curves saturate at higher prediction MAE values when smaller values of N_{cap} are employed. Nonetheless, the figure demonstrates that if additional training data is available, then the prediction accuracy for MOB-ML with RC can be substantially improved while capping the size of the largest cluster.

Figure 10b demonstrates the actual effect of capping on the parallelized training time, plotting the prediction MAE versus parallelized training time as a function of the number of training molecules. For reference, the results obtained using RC/LR/RFC and RC/GPR/RFC without capping are reproduced from Fig. 8. As is necessary, the RC/GPR/RFC results obtained with capping exactly overlap those obtained without capping when the number of training molecules is not greater than N_{cap} . However, for each value of N_{cap} , a sharp drop in the prediction MAE is seen when the number of training molecules begins to exceed N_{cap} , demonstrating that prediction accuracy can be greatly improved with minimal increase in parallelized training time. For example, it

is seen that for RC/GPR/RFC with $N_{\text{cap}} = 100$, chemical accuracy can be reached with only 7.4 s of parallelized training, slightly less than even RC/LR/RFC. For small values of N_{cap} , this prediction MAE eventually levels-off versus the training time, since the RFC training step becomes the dominant contribution to the training time.

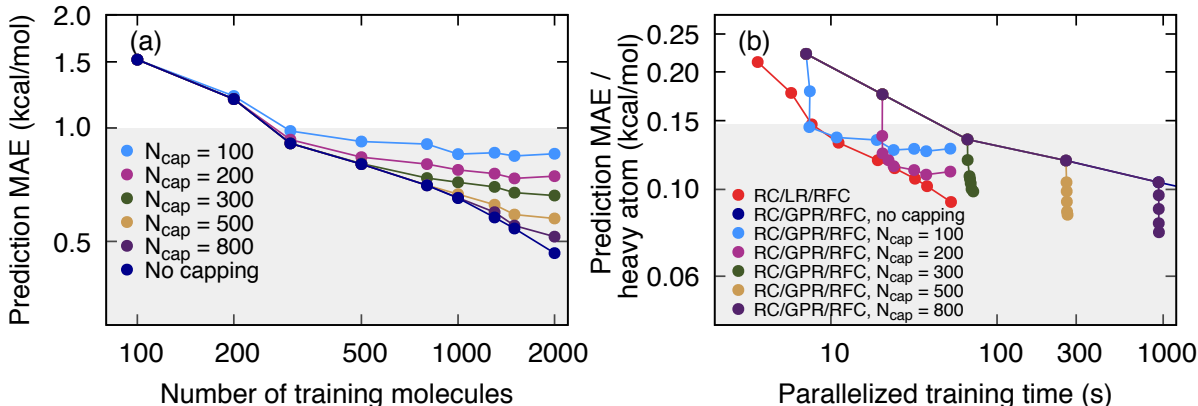


Figure 10: The effect of cluster-size capping on the prediction accuracy and training costs for MOB-ML with RC. Results reported for correlation energies at the MP2/cc-pVTZ level, with the training and test sets corresponding to non-overlapping subsets of the QM7b-T set of drug-like molecules with up to heavy seven atoms. (a) Prediction MAE versus the number of training molecules, with the clusters capped at various maximum sizes. The RC/GPR/RFC curve without capping is reproduced from Fig. 7a. (b) Prediction MAE per heavy atom versus parallelized training time as a function of the number of training molecules, as in Fig. 8. The results for MOB-ML with clustering and without capping cluster size (RC/LR/RFC, red; RC/GPR/RFC, blue) are reproduced from Fig. 8. Also, the results for RC/GPR/RFC with various capping sizes N_{cap} are shown. For part (a), the gray shaded area corresponds to a MAE of 1 kcal/mol, and for part (b), it corresponds to 1 kcal/mol per seven heavy atoms, to provide consistency with preceding figures. The prediction SEM is smaller than the plotted points.

5 Conclusions

Molecular-orbital-based (MOB) features offer a complete representation for mapping chemical space and a compact representation for evaluating correlation energies. In the current work, we take advantage of the intrinsic structure of MOB feature space, which cluster according to types of localized molecular orbitals, as well as the fact that orbital-pair contributions to the correlation energy contributions vary linearly with the MOB features, to overcome a fundamental bottleneck in

the efficiency of machine learning (ML) correlation energies. Specifically, we introduce a regression clustering (RC) approach in which MOB features and pair correlation energies are clustered according to their local linearity; we then individually regress these clusters and train a classifier for the prediction of cluster assignments on the basis of MOB features. This combined clustering/regression/classification approach is found to reduce MOB-ML training times by 3-4 orders of magnitude, while enabling prediction accuracies that are substantially improved over that which is possible using MOB-ML without clustering. The use of a random forest classifier for the cluster assignments, while better than alternatives that were explored, is found to be the limiting factor in terms of MOB-ML accuracy within this new approach, motivating future work on improved classifiers. This work provides a useful step towards that development of accurate, transferable, and scalable quantum ML methods to describe ever-broader swathes of chemical space.

Acknowledgement

This work emerged from a CMS 273 class project at Caltech that also involved Dmitry Burov, Jialin Song, Ying Shi Teh, and Dr. Tamara Husch, as well as Professors Kaushik Bhattacharya and Richard Murray; we thank these individuals for their ideas and contributions. This work is supported by the US Air Force Office of Scientific Research (AFOSR) grant FA9550-17-1-0102. M.W. acknowledges a postdoctoral fellowship from the Resnick Sustainability Institute. N.B.K. is supported, in part, by the US National Science Foundation (NSF) grant DMS 1818977, the US Office of Naval Research (ONR) grant N00014-17-1-2079, and the US Army Research Office (ARO) grant W911NF-12-2-0022. Computational resources were provided by the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the DOE Office of Science under contract DE-AC02-05CH11231.

Supporting Information Available

Figures in the supplementary information show the effect of averaging over independently trained MOB-ML-models (Fig. S1), the sensitivity of the prediction accuracy to the RC convergence threshold (Fig. S2), learning curves for various implementations of MOB-ML plotted on a log-log scale (Fig. S3), and a detailed breakdown of the parallelized wall-clock timings (Fig. S4). Tables in the supporting information provide the numerical data for the plots appearing in the main text.

References

- (1) Lavecchia, A. Machine-learning approaches in drug discovery: Methods and applications. *Drug Discov. Today* **2015**, *20*, 318–331.
- (2) Gawehn, E.; Hiss, J. A.; Schneider, G. Deep learning in drug discovery. *Mol. Inform.* **2016**, *35*, 3–14.
- (3) Popova, M.; Isayev, O.; Tropsha, A. Deep reinforcement learning for de novo drug design. *Sci. Adv.* **2018**, *4*, eaap7885.
- (4) Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; Riley, P. Molecular graph convolutions: Moving beyond fingerprints. *J. Comput. Aided Mol. Des.* **2016**, *30*, 595.
- (5) Mater, A. C.; Coote, M. L. Deep learning in chemistry. *J. Chem. Inf. Model* **2019**, *59*, 2545–2559.
- (6) Kim, E.; Huang, K.; Jegelka, S.; Olivetti, E. Virtual screening of inorganic materials synthesis parameters with deep learning. *npj Comput. Mater.* **2017**, *3*, 53.
- (7) Ren, F.; Ward, L.; Williams, T.; Laws, K. J.; Wolverton, C.; Hattrick-Simpers, J.; Mehta, A. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Sci. Adv.* **2018**, *4*, eaaq1566.

- (8) Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; Walsh, A. Machine learning for molecular and materials science. *Nature* **2018**, 559, 547–555.
- (9) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **2018**, 361, 360–365.
- (10) Wei, J. N.; Duvenaud, D.; Aspuru-Guzik, A. Neural networks for the prediction of organic chemistry reactions. *ACS Cent. Sci.* **2016**, 2, 725–732.
- (11) Raccuglia, P.; Elbert, K. C.; Adler, P. D. F.; Falk, C.; Wenny, M. B.; Mollo, A.; Zeller, M.; Friedler, S. A.; Schrier, J.; Norquist, A. J. Machine-learning-assisted materials discovery using failed experiments. *Nature* **2016**, 533, 73–76.
- (12) Ulissi, Z. W.; Medford, A. J.; Bligaard, T.; Nørskov, J. K. To address surface reaction network complexity using scaling relations machine learning and DFT calculations. *Nat. Commun.* **2017**, 8, 14621.
- (13) Segler, M. H. S.; Waller, M. P. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chem. - A Eur. J.* **2017**, 23, 5966–5971.
- (14) Segler, M. H. S.; Preuss, M.; Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* **2018**, 555, 604–610.
- (15) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **2017**, 8, 3192–3203.
- (16) Smith, J. S.; Nebgen, B. T.; Zubatyuk, R.; Lubbers, N.; Devereux, C.; Barros, K.; Tretiak, S.; Isayev, O.; Roitberg, A. E. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **2019**, 10, 2903.
- (17) Lubbers, N.; Smith, J. S.; Barros, K. Hierarchical modeling of molecular energies using a deep neural network. *J. Chem. Phys.* **2018**, 148, 241715.

- (18) Bartók, A. P.; Payne, M. C.; Kondor, R.; Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **2010**, *104*, 136403.
- (19) Rupp, M.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **2012**, *108*, 58301.
- (20) Montavon, G.; Rupp, M.; Gobre, V.; Vazquez-Mayagoitia, A.; Hansen, K.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Machine learning of molecular electronic properties in chemical compound space. *New J. Phys.* **2013**, *15*, 95003.
- (21) Hansen, K.; Montavon, G.; Biegler, F.; Fazli, S.; Rupp, M.; Scheffler, M.; von Lilienfeld, O. A.; Tkatchenko, A.; Müller, K.-R. Assessment and validation of machine learning methods for predicting molecular atomization energies. *J. Chem. Theory Comput.* **2013**, *9*, 3404.
- (22) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Big data meets quantum chemistry approximations: the Δ -machine learning approach. *J. Chem. Theory Comput.* **2015**, *11*, 2087.
- (23) Behler, J. Perspective: Machine learning potentials for atomistic simulations. *J. Chem. Phys.* **2016**, *145*, 170901.
- (24) Paesani, F. Getting the right answers for the right reasons: toward predictive molecular simulations of water with many-body potential energy functions. *Acc. Chem. Res.* **2016**, *49*, 1844.
- (25) Schütt, K. T.; Arbabzadah, F.; Chmiela, S.; Müller, K.-R.; Tkatchenko, A. Quantum-chemical insights from deep tensor neural networks. *Nat. Commun.* **2017**, *8*, 13890.
- (26) Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; Pande, V. MoleculeNet: a benchmark for molecular machine learning. *Chem. Sci.* **2018**, *9*, 513.

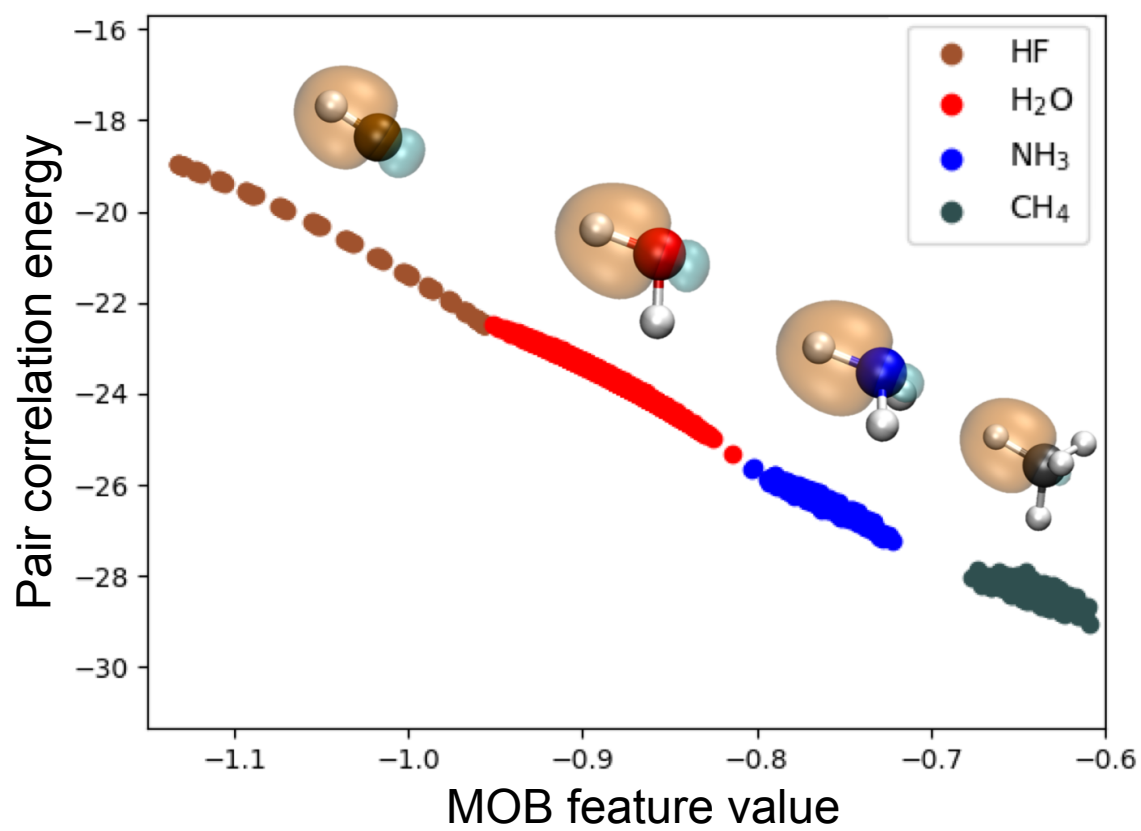
- (27) Nguyen, T. T.; Székely, E.; Imbalzano, G.; Behler, J.; Csányi, G.; Ceriotti, M.; Götz, A. W.; Paesani, F. Comparison of permutationally invariant polynomials, neural networks, and Gaussian approximation potentials in representing water interactions through many-body expansions. *J. Chem. Phys.* **2018**, *148*, 241725.
- (28) Fujikake, S.; Deringer, V. L.; Lee, T. H.; Krynski, M.; Elliott, S. R.; Csányi, G. Gaussian approximation potential modeling of lithium intercalation in carbon nanostructures. *J. Chem. Phys.* **2018**, *148*, 241714.
- (29) Li, H.; Collins, C.; Tanha, M.; Gordon, G. J.; Yaron, D. J. A density functional tight binding layer for deep learning of chemical Hamiltonians. *J. Chem. Theory Comput.* **2018**, *14*, 5764–5776.
- (30) Zhang, L.; Han, J.; Wang, H.; Car, R.; E, W. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **2018**, *120*, 143001.
- (31) Nandy, A.; Duan, C.; Janet, J. P.; Gugler, S.; Kulik, H. J. Strategies and software for machine learning accelerated discovery in transition metal chemistry. *Ind. Eng. Chem. Res.* **2018**, *57*, 13973–13986.
- (32) Grisafi, A.; Wilkins, D. M.; Willatt, M. J.; Ceriotti, M. Atomic-scale representation and statistical learning of tensorial properties. 2019, arXiv:1904.01623. arXiv.org ePrint archive. <https://arxiv.org/abs/1904.01623> (accessed Sep 04, 2019).
- (33) Grisafi, A.; Fabrizio, A.; Meyer, B.; Wilkins, D. M.; Corminboeuf, C.; Ceriotti, M. Transferable machine-learning model of the electron density. *ACS Cent. Sci.* **2019**, *5*, 57–64.
- (34) Bogojeski, M.; Brockherde, F.; Vogt-Maranto, L.; Li, L.; Tuckerman, M. E.; Burke, K.; Müller, K.-R. Efficient prediction of 3D electron densities using machine learning. 2018, arXiv:1811.06255. arXiv.org ePrint archive. <https://arxiv.org/abs/1811.06255> (accessed Sep 04, 2019).

- (35) Pereira, F.; Aires-de Sousa, J. Machine learning for the prediction of molecular dipole moments obtained by density functional theory. *J. Cheminf.* **2018**, *10*, 43.
- (36) Smith, J. S.; Nebgen, B. T.; Zubatyuk, R.; Lubbers, N.; Devereux, C.; Barros, K.; Tretiak, S.; Isayev, O.; Roitberg, A. E. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **2019**, *10*, 2903.
- (37) Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; von Lilienfeld, O. A. Electronic spectra from TDDFT and machine learning in chemical space. *J. Chem. Phys.* **2015**, *143*, 084111.
- (38) Gastegger, M.; Behler, J.; Marquetand, P. Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.* **2017**, *8*, 6924–6935.
- (39) Yao, K.; Herr, J. E.; Toth, D. W.; McKintyre, R.; Parkhill, J. The TensorMol-0.1 model chemistry: a neural network augmented with long-range physics. *Chem. Sci.* **2018**, *9*, 2261–2269.
- (40) Christensen, A. S.; Faber, F. A.; von Lilienfeld, O. A. Operators in quantum machine learning: Response properties in chemical space. *J. Chem. Phys.* **2019**, *150*, 064105.
- (41) Ghosh, K.; Stuke, A.; Todorovi, M.; Jrgensen, P. B.; Schmidt, M. N.; Vehtari, A.; Rinke, P. Deep learning spectroscopy: Neural networks for molecular excitation spectra. *Adv. Sci.* **2019**, *6*, 1801367.
- (42) Brockherde, F.; Vogt, L.; Li, L.; Tuckerman, M. E.; Burke, K.; Müller, K.-R. Bypassing the Kohn-Sham equations with machine learning. *Nat. Commun.* **2017**, *8*, 872.
- (43) McGibbon, R. T.; Taube, A. G.; Donchev, A. G.; Siva, K.; Hernández, F.; Hargus, C.; Law, K.-H.; Klepeis, J. L.; Shaw, D. E. Improving the accuracy of Møller-Plesset perturbation theory with neural networks. *J. Chem. Phys.* **2017**, *147*, 161725.
- (44) Nuddejima, T.; Ikabata, Y.; Seino, J.; Yoshikawa, T.; Nakai, H. Machine-learned electron

- correlation model based on correlation energy density at complete basis set limit. *J. Chem. Phys.* **2019**, *151*, 024104.
- (45) Townsend, J.; Vogiatzis, K. D. Data-Driven acceleration of the coupled-cluster singles and doubles iterative solver. *J. Phys. Chem. Lett.* **2019**, *10*, 4129–4135.
- (46) Sinitskiy, A. V.; Pande, V. S. Physical machine learning outperforms “human learning” in quantum chemistry. 2019, arXiv:1908.00971. arXiv.org ePrint archive. <https://arxiv.org/abs/1908.00971> (accessed Sep 04, 2019).
- (47) Welborn, M.; Cheng, L.; Miller, T. F., III Transferability in machine learning for electronic structure via the molecular orbital basis. *J. Chem. Theory Comput.* **2018**, *14*, 4772–4779.
- (48) Cheng, L.; Welborn, M.; Christensen, A. S.; Miller, T. F., III A universal density matrix functional from molecular orbital-based machine learning: Transferability across organic molecules. *J. Chem. Phys.* **2019**, *150*, 131103.
- (49) Møller, C.; Plesset, M. S. Note on an approximation treatment for many-electron systems. *Phys. Rev.* **1934**, *46*, 618.
- (50) Saebo, S.; Pulay, P. Local treatment of electron correlation. *Annu. Rev. Phys. Chem.* **1993**, *44*, 213–236.
- (51) Bartlett, R. J.; Watts, J. D.; Kucharski, S. A.; Noga, J. Non-iterative fifth-order triple and quadruple excitation energy corrections in correlated methods. *Chem. Phys. Lett.* **1990**, *165*, 513–522.
- (52) Schütz, M. Low-order scaling local electron correlation methods. III. Linear scaling local perturbative triples correction (T). *J. Chem. Phys.* **2000**, *113*, 9986–10001.
- (53) Nesbet, R. K. Brueckner’s theory and the method of superposition of configurations. *Phys. Rev.* **1958**, *109*, 1632.

- (54) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry*; Dover: Mineola, 1996; pp 231–239.
- (55) Späth, H. Correction to algorithm 39: clusterwise linear regression. *Computing* **1979**, *26*, 275.
- (56) Späth, H. Algorithm 39: clusterwise linear regression. *Computing* **1979**, 367–373.
- (57) Lloyd, S. Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **1982**, *28*, 129–137.
- (58) Montavon, G.; Rupp, M.; Gobre, V.; Vazquez-Mayagoitia, A.; Hansen, K.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Machine learning of molecular electronic properties in chemical compound space. *New J. Phys.* **2013**, *15*, 095003.
- (59) Blum, L. C.; Raymond, J.-L. 970 Million druglike small molecules for virtual screening in the chemical universe database GDB-13. *J. Am. Chem. Soc.* **2009**, *131*, 8732.
- (60) 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.
- (61) Pedregosa, F. et al. Scikit-learn: machine learning in python (v0.21.2). *J. Mach. Learn. Res.* **2011**, *12*, 2825.
- (62) GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>, since 2012.
- (63) Rasmussen, C. E.; Williams, C. K. I. *Gaussian processes for machine learning*; MIT Press: Cambridge, MA, 2006.
- (64) Criminisi, A.; Shotton, J.; Konukoglu, E. Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Found. Trends Comput. Graph. Vis.* **2012**, *7*, 81–227.
- (65) Fan, R.-E.; Chang, K.-W.; Hsieh, C.-J.; Wang, X.-R.; Lin, C.-J. LIBLINEAR: A library for large linear classification. *J. Mach. Learn. Res.* **2008**, *9*, 1871–1874.

- (66) Hastie, T.; Rosset, S.; Zhu, J.; Zou, H. Multi-class adaboost. *Stat. Its Interface* **2009**, 2, 349–360.
- (67) Cortes, C.; Jackel, L. D.; Solla, S. A.; Vapnik, V.; Denker, J. S. In *Advances in Neural Information Processing Systems 6*; Cowan, J. D., Tesauro, G., Alspector, J., Eds.; Morgan-Kaufmann, 1994; pp 327–334.
- (68) Baraldi, A.; Blonda, P. A survey of fuzzy clustering algorithms for pattern recognition. I. *IEEE Trans. Syst. Man Cybern. B Cybern.* **1999**, 29, 778–785.
- (69) Browning, N. J.; Ramakrishnan, R.; von Lilienfeld, O. A.; Roethlisberger, U. Genetic Optimization of Training Sets for Improved Machine Learning Models of Molecular Properties. *J. Phys. Chem. Lett.* **2017**, 8, 1351–1359.
- (70) Murphy, K. P. *Machine learning: a probabilistic perspective*; MIT Press: Cambridge, MA, 2012; pp 492–493.
- (71) Christensen, A. S.; Bratholm, L. A.; Faber, F. A.; Glowacki, D. R.; von Lilienfeld, O. A. FCHL revisited: faster and more accurate quantum machine learning. 2018, arXiv:1909.01946. arXiv.org ePrint archive. <https://arxiv.org/abs/1909.01946> (accessed Oct 18, 2019).
- (72) Faber, F. A.; Christensen, A. S.; Huang, B.; von Lilienfeld, O. A. Alchemical and structural distribution based representation for universal quantum machine learning. *J. Chem. Phys.* **2018**, 148, 241717.



TOC graphic