

A universal density matrix functional from molecular orbital-based machine learning: Transferability across organic molecules

Cite as: J. Chem. Phys. **150**, 131103 (2019); <https://doi.org/10.1063/1.5088393>

Submitted: 10 January 2019 . Accepted: 19 March 2019 . Published Online: 04 April 2019

Lixue Cheng , Matthew Welborn, Anders S. Christensen , and Thomas F. Miller 



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Cluster perturbation theory. I. Theoretical foundation for a coupled cluster target state and ground-state energies](#)

The Journal of Chemical Physics **150**, 134108 (2019); <https://doi.org/10.1063/1.5004037>

[Perspective: Computational chemistry software and its advancement as illustrated through three grand challenge cases for molecular science](#)

The Journal of Chemical Physics **149**, 180901 (2018); <https://doi.org/10.1063/1.5052551>

[Data sampling scheme for reproducing energies along reaction coordinates in high-dimensional neural network potentials](#)

The Journal of Chemical Physics **150**, 134103 (2019); <https://doi.org/10.1063/1.5078394>

The Journal
of Chemical Physics

2018 EDITORS' CHOICE

READ NOW!



A universal density matrix functional from molecular orbital-based machine learning: Transferability across organic molecules

Cite as: J. Chem. Phys. 150, 131103 (2019); doi: 10.1063/1.5088393

Submitted: 10 January 2019 • Accepted: 19 March 2019 •

Published Online: 4 April 2019



View Online



Export Citation



CrossMark

Lixue Cheng,¹  Matthew Welborn,¹ Anders S. Christensen,²  and Thomas F. Miller III^{1,a)} 

AFFILIATIONS

¹Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, California 91125, USA

²Institute of Physical Chemistry and National Center for Computational Design and Discovery of Novel Materials, Department of Chemistry, University of Basel, Basel, Switzerland

^{a)}Electronic mail: tfm@caltech.edu.

ABSTRACT

We address the degree to which machine learning (ML) can be used to accurately and transferably predict post-Hartree-Fock correlation energies. Refined strategies for feature design and selection are presented, and the molecular-orbital-based machine learning (MOB-ML) method is applied to several test systems. Strikingly, for the second-order Møller-Plessett perturbation theory, coupled cluster with singles and doubles (CCSD), and CCSD with perturbative triples levels of theory, it is shown that the thermally accessible (350 K) potential energy surface for a single water molecule can be described to within 1 mhartree using a model that is trained from only a single reference calculation at a randomized geometry. To explore the breadth of chemical diversity that can be described, MOB-ML is also applied to a new dataset of thermalized (350 K) geometries of 7211 organic models with up to seven heavy atoms. In comparison with the previously reported Δ -ML method, MOB-ML is shown to reach chemical accuracy with threefold fewer training geometries. Finally, a transferability test in which models trained for seven-heavy-atom systems are used to predict energies for thirteen-heavy-atom systems reveals that MOB-ML reaches chemical accuracy with 36-fold fewer training calculations than Δ -ML (140 vs 5000 training calculations).

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5088393>

I. INTRODUCTION

Machine learning (ML) has recently seen wide application in chemistry, including the fields of drug discovery,^{1–3} materials design,^{4–7} and reaction prediction.^{8–12} In the context of quantum chemistry, much work has focused on predicting electronic energies or densities based on atom- or geometry-specific features,^{13–33} although other strategies have also been employed.³⁴ Recently, we reported an accurate and transferable molecular-orbital-based machine learning (MOB-ML) approach to the prediction of correlated wavefunction energies based on input features from a self-consistent field calculation such as the Hartree-Fock (HF) method.³⁵ In this communication, we present refinements to the MOB-ML method with comparisons to test cases from our previous work. We then demonstrate the performance of MOB-ML across a broad swath of chemical space, as represented by the QM7b³⁶ and GDB-13³⁷ test sets of organic molecules.

II. THEORY

The current work aims to predict post-Hartree-Fock correlated wavefunction energies using features of the Hartree-Fock molecular orbitals (MOs). The starting point for the MOB-ML method³⁵ is that the correlation energy can be decomposed into pairwise occupied MO contributions^{38,39}

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

where the pair correlation energy ε_{ij} can be written as a functional of the full set of MOs, $\{\phi_p\}$, appropriately indexed by i and j

$$\varepsilon_{ij} = \varepsilon[\{\phi_p\}^{ij}]. \quad (2)$$

The functional ε is universal across all chemical systems; for a given level of correlated wavefunction theory, there is a

corresponding ε that maps the HF MOs to the pair correlation energy, regardless of the molecular composition or geometry. Furthermore, ε simultaneously describes the pair correlation energy for all pairs of occupied MOs (i.e., the functional form of ε does not depend on i and j). For example, in second-order Møller-Plessett perturbation theory (MP2),⁴⁰ the pair correlation energies are

$$\varepsilon_{ij}^{\text{MP2}} = \frac{1}{4} \sum_{ab}^{\text{virt}} \frac{|(ij||ab)|^2}{e_a + e_b - e_i - e_j}, \quad (3)$$

where a and b index virtual MOs, e_p is the Hartree-Fock orbital energy corresponding to MO ϕ_p , and $\langle ij||ab \rangle$ are antisymmetrized electron repulsion integrals.³⁹ A corresponding expression for the pair correlation energy exists for any post-Hartree-Fock method, but it is typically costly to evaluate in closed form.

In MOB-ML, a machine learning model is constructed for the pair energy functional

$$\varepsilon_{ij} \approx \varepsilon^{\text{ML}}[\mathbf{f}_{ij}], \quad (4)$$

where \mathbf{f}_{ij} denotes a vector of features associated with MOs i and j . Equation (4) thus presents the opportunity for the machine learning of a universal density matrix functional for correlated wavefunction energies, which can be evaluated at the cost of the MO calculation.

Following our previous work,³⁵ the features \mathbf{f}_{ij} correspond to unique elements of the Fock (**F**), Coulomb (**J**), and exchange (**K**) matrices between ϕ_i , ϕ_j , and the set of virtual orbitals. In the current work, we additionally include features associated with matrix elements between pairs of occupied orbitals for which one member of the pair differs from ϕ_i or ϕ_j (i.e., non- i,j occupied MO pairs). The feature vector takes the form

$$\mathbf{f}_{ij} = (F_{ii}, F_{ij}, F_{jj}, \mathbf{F}_i^o, \mathbf{F}_j^o, \mathbf{F}_{ij}^{ov}, J_{ii}, J_{ij}, J_{jj}, \mathbf{J}_i^o, \mathbf{J}_j^o, \mathbf{J}_{ij}^v, \mathbf{J}_i^{vv}, \mathbf{J}_j^{vv}, \mathbf{J}_{ij}^{vv}, K_{ij}, \mathbf{K}_i^o, \mathbf{K}_j^o, \mathbf{K}_i^v, \mathbf{K}_j^v, \mathbf{K}_{ij}^{vv}), \quad (5)$$

where for a given matrix (**F**, **J**, or **K**), the superscript o denotes a row of its occupied-occupied block, the superscript v denotes a row of its occupied-virtual block, and the superscript vv denotes its virtual-virtual block. Redundant elements are removed such that the virtual-virtual block is represented by its upper triangle, and the diagonal elements of **K** (which are identical to those of **J**) are omitted. To increase transferability and accuracy, we choose ϕ_i and ϕ_j to be localized molecular orbitals (LMOs) rather than canonical MOs and employ valence virtual LMOs⁴¹ in place of the set of all virtual MOs (as detailed in Ref. 35). We separate Eq. (4) to independently machine learn the cases of $i = j$ and $i \neq j$

$$\varepsilon_{ij} \approx \begin{cases} \varepsilon_d^{\text{ML}}[\mathbf{f}_i] & \text{if } i = j \\ \varepsilon_o^{\text{ML}}[\mathbf{f}_{ij}] & \text{if } i \neq j, \end{cases} \quad (6)$$

where \mathbf{f}_i denotes \mathbf{f}_{ii} [Eq. (5)] with redundant elements removed; by separating the pair energies in this way, we avoid the situation where a single ML model is required to distinguish between the cases of $i = j$ and ϕ_i being nearly degenerate to ϕ_j , a distinction which can represent a sharp variation in the function to be learned.

In the current work, several technical refinements are introduced to improve training efficiency (i.e., the accuracy and

transferability of the model as a function of the number of training examples). These are now described.

A. Occupied LMO symmetrization

The feature vector is preprocessed to specify a canonical ordering of the occupied and valence virtual LMO pairs. This reduces permutation of elements in the feature vector, resulting in greater ML training efficiency. Matrix elements M_{ij} (**M** = **F**, **J**, **K**) associated with ϕ_i and ϕ_j are rotated into gerade and ungerade combinations

$$\begin{aligned} M_{ii} &\leftarrow \frac{1}{2}M_{ii} + \frac{1}{2}M_{jj} + M_{ij}, \\ M_{jj} &\leftarrow \frac{1}{2}M_{ii} + \frac{1}{2}M_{jj} - M_{ij}, \\ M_{ij} &\leftarrow \frac{1}{2}M_{ii} - \frac{1}{2}M_{jj}, \\ M_{ip} &\leftarrow \frac{1}{\sqrt{2}}M_{ip} + \frac{1}{\sqrt{2}}M_{jp}, \\ M_{jp} &\leftarrow \frac{1}{\sqrt{2}}M_{ip} - \frac{1}{\sqrt{2}}M_{jp}, \end{aligned} \quad (7)$$

with the sign convention that F_{ij} is negative. Here, p indexes any LMO other than i or j (i.e., an occupied LMO k , such that $i \neq k \neq j$, or a valence virtual LMO).

B. LMO sorting

The valence virtual LMO pairs are sorted by increasing distance from the occupied orbitals ϕ_i and ϕ_j . Sorting in this way ensures that features corresponding to valence virtual LMOs are listed in decreasing order of heuristic importance and that the mapping between valence virtual LMOs and their associated features is roughly preserved. We recognize this issue could also potentially be addressed through the use of symmetry functions,⁴² but these are not employed in the current work.

For purposes of sorting, distance is defined as

$$R_a^{ij} = \|\langle \phi_i | \hat{R} | \phi_i \rangle - \langle \phi_a | \hat{R} | \phi_a \rangle\| + \|\langle \phi_j | \hat{R} | \phi_j \rangle - \langle \phi_a | \hat{R} | \phi_a \rangle\|, \quad (8)$$

where ϕ_a is a valence virtual LMO, \hat{R} is the Cartesian position operator, and $\|\cdot\|$ denotes the 2-norm. $\|\langle \phi_i | \hat{R} | \phi_i \rangle - \langle \phi_a | \hat{R} | \phi_a \rangle\|$ represents the Euclidean distance between the centroids of orbital i and orbital a . Previously,³⁵ distances were defined based on Coulomb repulsion, which was found to sometimes lead to inconsistent sorting in systems with strongly polarized bonds. The non- i, j occupied LMO pairs are sorted in the same manner as the valence virtual LMO pairs.

C. Orbital localization

We employ Boys localization⁴³ to obtain the occupied LMOs, rather than intrinsic bond orbital (IBO) localization⁴¹ which was employed in our previous work.³⁵ Particularly, for molecules that include triple bonds or multiple lone pairs, it is found that Boys localization provides more consistent localization as a function of small geometry changes than IBO localization, and the chemically unintuitive mixing of σ and π bonds in Boys localization (“banana bonds”)⁴⁴ does not present a problem for the MOB-ML method.

D. Feature selection

Prior to training, automatic feature selection is performed using random forest regression⁴⁵ with the mean decrease of accuracy criterion (sometimes called permutation importance).⁴⁶ This technique was found to be more effective than our previous use³⁵ of the Gini importance score⁴⁵ which led to worse accuracy and failed to select any features for the case of methane.

The reason for using feature selection in this way is twofold. First, Gaussian process regression (GPR) performance is known to degrade for high-dimensional datasets (in practice, 50–100 features),⁴⁷ and second, the use of the full feature set with small molecules can lead to overfitting as features can become correlated.

A reference PYTHON implementation for generating MOB-ML features is provided online.⁴⁸

III. COMPUTATIONAL DETAILS

Results are presented for a single water molecule; a series of alkane molecules; a thermalized version of the QM7b set of 7211 molecules with up to seven C, O, N, S, and Cl heavy atoms; and a thermalized version of the GDB-13 set of molecules with thirteen C, O, N, S, and Cl heavy atoms. All datasets employed in this work are provided in the [supplementary material](#).

Training and test geometries are sampled at 50 fs intervals from *ab initio* molecular dynamics trajectories performed with the Q-CHEM 5.0 software package,⁴⁹ using the B3LYP^{50–53}/6-31g*⁵⁴ level of theory and a Langevin thermostat⁵⁵ at 350 K.

The features and training pair energies associated with these geometries are computed using the MOLPRO 2018.0 software package⁵⁶ in a cc-pVTZ basis set unless otherwise noted.⁵⁷ Valence virtual orbitals used in feature construction are determined with the IBO method.⁴¹ Reference pair correlation energies are computed with second-order Møller-Plessett perturbation theory (MP2)^{40,58} and coupled cluster with singles and doubles (CCSD)^{59,60} as well as with perturbative triples [CCSD(T)].^{61,62} Density fitting for both Coulomb and exchange integrals⁶³ is employed for all results below except those corresponding to the water molecule. The frozen core approximation is used in all cases.

Gaussian process regression (GPR)⁶⁴ is employed to machine learn ϵ_d^{ML} and ϵ_o^{ML} [Eq. (6)] using the GPY 1.9.6 software package.⁶⁵ The GPR kernel is Matérn 5/2 with white noise regularization.⁶⁴ Kernel hyperparameters are optimized with respect to the log marginal likelihood objective for the water and alkane series results, as well as for ϵ_d^{ML} of the QM7b results. We use the Matérn 3/2 kernel instead of the Matérn 5/2 kernel for the case of ϵ_o^{ML} for QM7b results as it was empirically found to yield slightly better accuracy.⁶⁶

Feature selection is performed using the random forest regression implementation in the SCIKIT-LEARN v0.20.0 package.⁶⁷

IV. RESULTS

The ML model of Eq. (6) is a universal functional for any molecular Hamiltonian. In principle, with an adequate feature list and unlimited training data (and time), it should accurately and simultaneously describe all molecular systems. In practice, we must train the ML model using a truncated feature list and finite data. These choices determine the accuracy of the model.

Below, we examine the performance of the MOB-ML method in three increasingly broad regions of chemical space: (i) training on randomized water molecule geometries and predicting the energies of other water molecule geometries, (ii) training on geometries of short alkanes and predicting the energies of longer alkanes, and (iii) training on a small set of organic molecules and predicting the energies of a broader set of organic molecules. The first two test cases were introduced in our previous work,³⁵ and we explore how the refined methodology reported herein leads to improvements in accuracy and transferability. The last case represents a demanding new test of transferability across chemical space. In all cases, we report the ML prediction accuracy as a function of the number of training examples.

As a first example, we consider the performance of MOB-ML for a single water molecule. A separate model is trained to predict the correlation energy at the MP2, CCSD, and CCSD(T) levels of theory, using reference calculations on a subset of 1000 randomized water geometries to predict the correlation energy for the remainder. Feature selection with an importance threshold of 1.00×10^{-3} results in 12, 11, and 10 features for ϵ_o^{ML} for MP2, CCSD, and CCSD(T), respectively; ten features are selected for ϵ_d^{ML} for all three post-Hartree-Fock methods.

Figure 1 presents the test set prediction accuracy of each MOB-ML model as a function of the number of training geometries (i.e., the “learning curve”). MOB-ML predictions are shown for MP2, CCSD, and CCSD(T), and the model shows the same level of accuracy for all three methods. Remarkably, all three models achieve a prediction mean absolute error (MAE) of 1 mhartree when trained on only a single water geometry, indicating that only a single reference calculation is needed to provide chemical accuracy for the remaining 999 geometries at each level of theory. Since it contains 10 distinct LMO pairs, this single geometry provides enough information to yield a chemically accurate MOB-ML model for the global thermally accessible potential energy surface.

For all three methods (Fig. 1), the learning curve exhibits the expected⁶⁸ power-law behavior as a function of training data, and the total error reaches microhartree accuracy with tens of water training geometries. As compared to our previous results, where

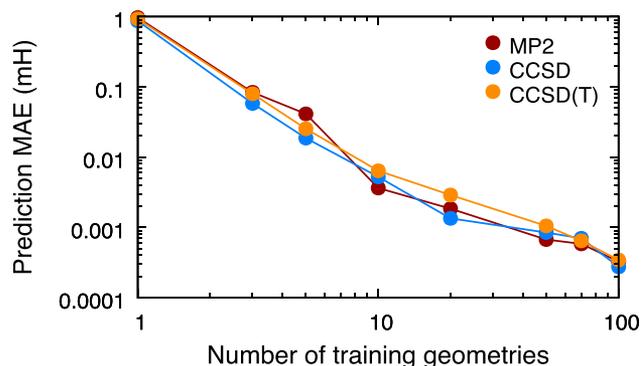


FIG. 1. Learning curves for MOB-ML models trained on the water molecule and used to predict the correlation energy of different water molecule geometries at three levels of post-Hartree-Fock theory. Prediction errors are summarized in terms of mean absolute error (MAE).

training on 200 geometries resulted in a prediction MAE of 0.027 mhartree for the case of CCSD,³⁵ the current implementation of the MOB-ML model is substantially improved; the improvement for this case stems primarily from the use of Boys localization,⁴³ which specifies unique and consistent LMOs corresponding to the oxygen lone pairs.

Next, we explore the transferability of MOB-ML predictions for a model that is trained on thermalized geometries of short alkanes and then used for predictions on thermalized geometries of larger and more branched alkanes (n-butane and isobutane). For these predictions, the absolute zero of energy is shifted for each molecule to compare relative energies on its potential energy surface (i.e., parallelity errors are removed). These shifts are reported in the caption of Fig. 2; for no other results reported in the paper are parallelity errors removed.

In our previous work,³⁵ this test was performed using training data that combined 100 geometries of methane, 300 of ethane, and 50 of propane; the resulting predictions are reproduced in Fig. 2(a). This earlier implementation of MOB-ML led to predictions for

n-butane and isobutane with substantial errors (0.59 mhartree for n-butane and 0.93 mhartree for isobutane) and a noticeable skew with respect to the true correlation energy.

The predictions of MOB-ML in the current work [Fig. 2(b)] are markedly improved. First, the overall prediction accuracy is improved for all four summary statistics (inset in Fig. 2) despite substantial reduction in the number of training examples used. (The current work uses only 50 geometries of ethane, 20 geometries of propane, and no methane data.) Second, n-butane and isobutane are predicted with nearly identical accuracy. Finally, the prediction errors are no longer skewed as a function of true correlation energy. The primary methodological sources of these improvements are found to be symmetrization of occupied orbitals [Eq. (7)] and the improved feature selection methodology. The MOB-ML features in the current work are selected with an importance threshold of 1×10^{-4} , resulting in 27 features for ϵ_d^{ML} and 12 features for ϵ_o^{ML} ; the results presented in Fig. 2(b) for CCSD(T) are qualitatively identical to those obtained for CCSD (not shown).

We now examine the transferability of the MOB-ML method across a broad swath of chemical space. Specifically, we consider the QM7b dataset,³⁶ which is comprised of 7211 plausible organic molecules with 7 or fewer heavy atoms. The chemical elements in QM7b are limited to those likely to be found in drug-like compounds: C, H, O, N, S, and Cl. We refer to the dataset used herein as QM7b-T to reflect the fact that it contains geometries sampled at a temperature of 350 K (as described in Sec. II), as opposed to density functional theory optimized geometries. The MOB-ML model is trained on a randomly chosen subset of QM7b-T molecules and used to predict the correlation energy of the remainder. Active learning was also tested as a training data selection strategy, but was not found to improve the predictions in the regime of chemical accuracy, and in fact led to slightly worse transferability.

For comparison, a Δ -ML model²¹ was trained on the same molecules using kernel-ridge regression using the Faber-Christensen-Huang-Lilienfeld (FHCL) representation⁷⁰ with a Gaussian kernel function (FCHL/ Δ -ML), as implemented in the QML package.⁷¹ All hyperparameters of the model were set to those obtained in Ref. 70, which have previously been demonstrated to work well for datasets containing structures similar to those in QM7b-T.⁷¹

A possible source of concern for MOB-ML is that the number of selected features would grow with the chemical complexity of the training data. For example, 27 features for ϵ_d^{ML} and 12 features for ϵ_o^{ML} were selected in the alkane test case using the ethane + propane training data [Fig. 2(b)], whereas only 10 features for ϵ_d^{ML} and 10 features for ϵ_o^{ML} were selected for the water test case at the CCSD(T) level of theory (Fig. 1). To examine this, we perform feature selection on increasing numbers of randomly selected molecules from the QM7b-T dataset. Table I presents two statistics on the feature importance as a function of the number of training molecules: (i) the number of “important features” (i.e., those whose permutation importance⁴⁶ exceeds a set threshold of 2×10^{-4} and 5×10^{-5} for ϵ_d^{ML} and ϵ_o^{ML} , respectively) and (ii) the inverse participation ratio⁷² of the feature importance scores. The latter is a threshold-less measure of the number of important features; it takes a value of 1 when only 1 feature has nonzero importance and N when all N features have equal importance. Although the QM7b-T dataset contains many different chemical elements and bonding

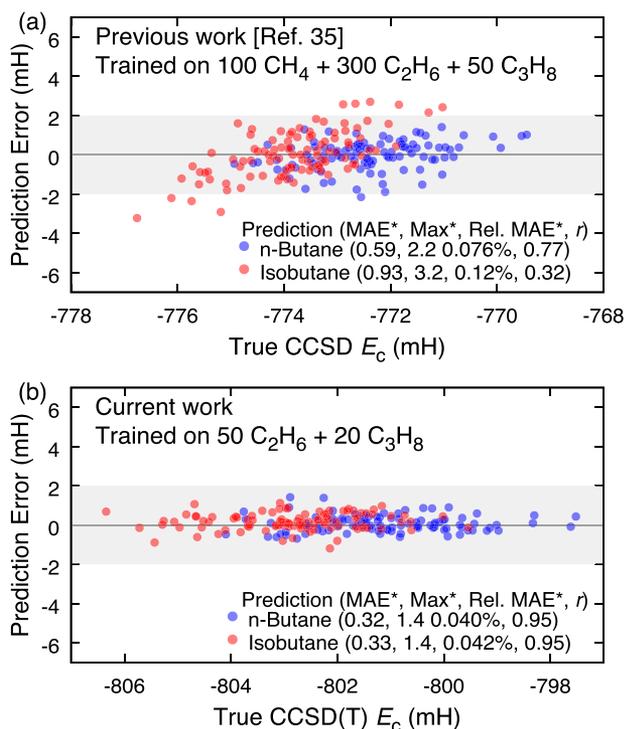


FIG. 2. MOB-ML predictions of the correlation energy for 100 n-butane and isobutane geometries, using MOB-ML features described in the (b) current work, compared to (a) the MOB-ML features of Ref. 35. Training sets are indicated in each panel of the figure. MOB-ML prediction errors are plotted vs the (a) true CCSD correlation energy and (b) true CCSD(T) correlation energy. To remove parallelity errors, a global shift is applied to the predictions of n-butane and isobutane by (a) 3.3 and 0.73 mhartree and (b) 0.90 and 0.17 mhartree respectively. Summary statistics that include this shift (indicated by an asterisk) are presented, consisting of mean absolute error (MAE*), maximum absolute error (Max*), MAE* as a percentage of E_c (Rel. MAE*), and Pearson correlation coefficient (r).⁶⁹ The gray shaded region corresponds to errors of ± 2 mhartree.

TABLE I. Number of features selected as a function of the number of randomly chosen training molecules for the QM7b-T dataset at the CCSD(T)/cc-pVDZ level. The number of features that exceeds an importance threshold as well as the inverse participation ratio (IPR) of the feature importance scores is reported (see text).

Training size	No. of important features		Feature weight IPR	
	ϵ_d^{ML}	ϵ_o^{ML}	ϵ_d^{ML}	ϵ_o^{ML}
20	50	28	4.720	1.116
50	46	28	3.718	1.097
100	46	26	3.450	1.115
200	42	24	3.430	1.120

motifs, Table I reveals that the selected features remain compact and do not grow with the number of training molecules. Indeed, for a large number of training molecules, the number of selected features slightly decreases, reaching 42 and 24 selected features for ϵ_d^{ML} and ϵ_o^{ML} , respectively.

The learning curves for MOB-ML models trained at the MP2/cc-pVTZ and CCSD(T)/cc-pVDZ levels of theory are shown in Fig. 3(a), as well as the FCHL/ Δ -ML learning curve for MP2/cc-pVTZ. At the MP2 level of theory, the MOB-ML model achieves an accuracy of 2 mhartree with 110 training calculations (representing 1.5% of the molecules in the QM7b-T dataset), whereas FCHL/ Δ -ML requires over 300 training geometries to reach the same accuracy threshold. Fig. 3(a) also illustrates the relative insensitivity of MOB-ML to the level of electronic structure theory, with the learning curve for CCSD(T)/cc-pVDZ reaching 2 mhartree accuracy with 140 training calculations. An analysis of the sensitivity of the MOB-ML predictions to the number of selected features is presented in the supplementary material, Fig. S1, which indicates that the reported results are robust with respect to the number of selected features.

As a final test of transferability of the MOB-ML and FCHL/ Δ -ML methods across chemical space, Figs. 3(b) and 3(c) show results in which the ML methods are trained on QM7b-T molecules and then used to predict results for a dataset of 13-heavy-atom organic molecules at thermalized geometries, GDB-13-T, which includes six thermally sampled geometries each of 1000 13-heavy-atom organic molecules chosen randomly from the GDB-13 dataset.³⁷ Like QM7b, the members of GDB-13 contain C, H, N, O, S, and Cl. The size of these molecules precludes the use of coupled cluster theory to generate reference data; we therefore make comparison at the MP2/cc-pVTZ level of theory, noting that MOB-ML has consistently been shown to be insensitive to the employed post-Hartree-Fock method [as shown in Fig. 3(a)]. Transfer learning results as a function of the number of training molecules are presented in Fig. 3(b) (on a linear-linear scale) and Fig. 3(c) (on a log-log scale).

Using the MOB-ML model that is trained on 110 seven-heavy-atom molecules (corresponding to a prediction MAE of 1.89 mhartree for QM7b-T), we observe a prediction MAE of 3.88 mhartree for GDB-13-T. Expressed in terms of size-intensive quantities, the prediction MAE per heavy atom is 0.277 mhartree and 0.298 mhartree for QM7b-T and GDB-13-T, respectively, indicating that the accuracy of the MOB-ML results is only slightly worse

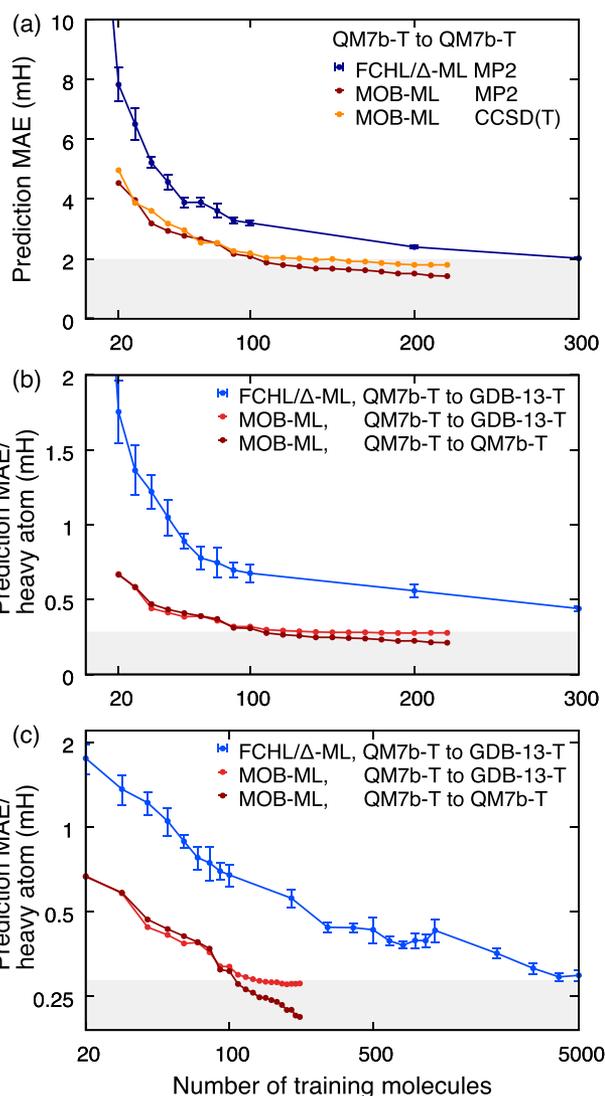


FIG. 3. Learning curves for MOB-ML trained on QM7b-T and applied to QM7b-T and GDB-13-T (see text for the definition of these datasets). FCHL/ Δ -ML⁷⁰ results are provided for comparison. (a) Predictions are made for QM7b-T at the MP2/cc-pVTZ (red) and CCSD(T)/cc-pVDZ (orange) levels of theory. (b) Using the same models trained on QM7b-T, predictions are made for GDB-13-T and reported in terms of MAE per heavy atom. (MOB-ML predictions for QM7b-T are included for reference.) (c) As in the previous panel but plotted on a logarithmic scale and extended to show the full range of FCHL/ Δ -ML predictions. Error bars for FCHL/ Δ -ML represent prediction standard errors of the mean as measured over 10 models. The gray shaded area corresponds to errors of 2 mhartree per 7 heavy atoms.

when the model is transferred to the dataset of larger molecules. On a per-heavy-atom basis, MOB-ML reaches chemical accuracy with the same number of QM7b-T training calculations (approximately 100), regardless of whether it is tested on QM7b-T or GDB-13-T.

In contrast with MOB-ML, the FCHL/ Δ -ML method is found to be significantly less transferable from QM7b-T to GDB-13-T. For

models trained using 100 seven-heavy-atom molecules, the MAE per heavy atom of FCHL/ Δ -ML is over twice that of MOB-ML [Fig. 3(b)]. Moreover, whereas MOB-ML reaches the per-heavy-atom chemical accuracy threshold with 140 training calculations, the FCHL/ Δ -ML method only reaches that threshold with 5000 training calculations.

V. CONCLUSIONS

Molecular-orbital-based machine learning (MOB-ML) has been shown to be a simple and strikingly accurate strategy for predicting correlated wavefunction energies at the cost of a Hartree-Fock calculation benefiting from the intrinsic transferability of the localized molecular orbital representation. The starting point for the MOB-ML method is a rigorous mapping from the Hartree-Fock molecular orbitals to the total correlation energy, which ensures that the use of sufficient training data and molecular orbital features will produce a model that matches the corresponding correlated wavefunction method across the entirety of chemical space. The current work explores this possibility within the subspace of organic molecules. It is shown that MOB-ML predicts energies of the QM7b-T dataset to within a 2 mhartree accuracy using only 110 training calculations at the MP2/cc-pVTZ level of theory and using 140 training calculations at the CCSD(T)/cc-pVDZ level of theory. Direct comparison with FCHL/ Δ -ML reveals that MOB-ML is three-fold more efficient in reaching chemical accuracy for describing QM7b-T. Furthermore, a transferability test of a MOB-ML model trained on QM7b-T to GDB-13-T reveals that MOB-ML exhibits negligible degradation in accuracy; as a result, chemical accuracy is achieved with 36-times fewer training calculations using MOB-ML vs FCHL/ Δ -ML. These results suggest that MOB-ML provides a promising approach toward the development of density matrix functionals that are applicable across broad swathes of chemical space.

SUPPLEMENTARY MATERIAL

The datasets used in this work are available for download;⁷³ they include MOB-ML features, HF energies, pair correlation energies, and geometries. MOB-ML and FCHL/ Δ -ML predictions corresponding to Fig. 3 and an analysis of the sensitivity of the results in Fig. 3 to the number of selected features are available in the [supplementary material](#). A reference PYTHON implementation for generating MOB-ML features is available at <https://github.com/thomasfmiller/MOB-ML>.

ACKNOWLEDGMENTS

We thank Daniel Smith (Molecular Sciences Software Institute) and Alberto Gobbi (Genentech) for a helpful discussion about available training datasets. T.F.M. acknowledges support from AFOSR Award No. FA9550-17-1-0102. A.S.C. acknowledges support from the National Centre of Competence in Research (NCCR) Materials Revolution: Computational Design and Discovery of Novel Materials (MARVEL) of the Swiss National Science Foundation (SNSF). We also acknowledge support from the Resnick Sustainability Institute postdoctoral fellowship (M.W.) and the Camille Dreyfus Teacher-Scholar Award (T.F.M.). 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 No. DE-AC02-05CH11231.

REFERENCES

- ¹A. Lavecchia, "Machine-learning approaches in drug discovery: Methods and applications," *Drug Discovery Today* **20**, 318–331 (2015).
- ²E. Gawehn, J. A. Hiss, and G. Schneider, "Deep learning in drug discovery," *Mol. Inf.* **35**, 3–14 (2016).
- ³M. Popova, O. Isayev, and A. Tropsha, "Deep reinforcement learning for de novo drug design," *Sci. Adv.* **4**, eaap7885 (2018).
- ⁴E. Kim, K. Huang, S. Jegelka, and E. Olivetti, "Virtual screening of inorganic materials synthesis parameters with deep learning," *npj Comput. Mater.* **3**, 53 (2017).
- ⁵F. Ren, L. Ward, T. Williams, K. J. Laws, C. Wolverton, J. Hatrick-Simpers, and A. Mehta, "Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments," *Sci. Adv.* **4**, eaq1566 (2018).
- ⁶K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, "Machine learning for molecular and materials science," *Nature* **559**, 547–555 (2018).
- ⁷B. Sanchez-Lengeling and A. Aspuru-Guzik, "Inverse molecular design using machine learning: Generative models for matter engineering," *Science* **361**, 360–365 (2018).
- ⁸J. N. Wei, D. Duvenaud, and A. Aspuru-Guzik, "Neural networks for the prediction of organic chemistry reactions," *ACS Cent. Sci.* **2**, 725–732 (2016).
- ⁹P. Raccuglia, K. C. Elbert, P. D. F. Adler, C. Falk, M. B. Wenny, A. Mollo, M. Zeller, S. A. Friedler, J. Schrier, and A. J. Norquist, "Machine-learning-assisted materials discovery using failed experiments," *Nature* **533**, 73–76 (2016).
- ¹⁰Z. W. Ulissi, A. J. Medford, T. Bligaard, and J. K. Nørskov, "To address surface reaction network complexity using scaling relations machine learning and DFT calculations," *Nat. Commun.* **8**, 14621 (2017).
- ¹¹M. H. S. Segler and M. P. Waller, "Neural-symbolic machine learning for retrosynthesis and reaction prediction," *Chem. - Eur. J.* **23**, 5966–5971 (2017).
- ¹²M. H. S. Segler, M. Preuss, and M. P. Waller, "Planning chemical syntheses with deep neural networks and symbolic AI," *Nature* **555**, 604–610 (2018).
- ¹³J. S. Smith, O. Isayev, and A. E. Roitberg, "ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost," *Chem. Sci.* **8**, 3192–3203 (2017).
- ¹⁴J. S. Smith, B. T. Nebgen, R. Zubatyuk, N. Lubbers, C. Devereux, K. Barros, S. Tretiak, O. Isayev, and A. Roitberg, "Outsmarting quantum chemistry through transfer learning," <https://doi.org/10.26434/chemrxiv.6744440.v1> (2018).
- ¹⁵N. Lubbers, J. S. Smith, and K. Barros, "Hierarchical modeling of molecular energies using a deep neural network," *J. Chem. Phys.* **148**, 241715 (2018).
- ¹⁶A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, "Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons," *Phys. Rev. Lett.* **104**, 136403 (2010).
- ¹⁷M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," *Phys. Rev. Lett.* **108**, 058301 (2012).
- ¹⁸D. M. Wilkins, A. Grisafi, Y. Yang, K. U. Lao, R. A. DiStasio Jr., and M. Ceriotti, "Accurate molecular polarizabilities with coupled cluster theory and machine learning," *Proc. Natl. Acad. Sci. U. S. A.* **116**, 3401–3406 (2019).
- ¹⁹K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, "Assessment and validation of machine learning methods for predicting molecular atomization energies," *J. Chem. Theory Comput.* **9**, 3404 (2013).
- ²⁰P. Gasparotto and M. Ceriotti, "Recognizing molecular patterns by machine learning: An agnostic structural definition of the hydrogen bond," *J. Chem. Phys.* **141**, 174110 (2014).
- ²¹R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, "Big data meets quantum chemistry approximations: The Δ -machine learning approach," *J. Chem. Theory Comput.* **11**, 2087 (2015).

- ²²J. Behler, "Perspective: Machine learning potentials for atomistic simulations," *J. Chem. Phys.* **145**, 170901 (2016).
- ²³S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley, "Molecular graph convolutions: Moving beyond fingerprints," *J. Comput.-Aided Mol. Des.* **30**, 595 (2016).
- ²⁴F. Paesani, "Getting the right answers for the right reasons: Toward predictive molecular simulations of water with many-body potential energy functions," *Acc. Chem. Res.* **49**, 1844 (2016).
- ²⁵K. T. Schütt, F. Arbabzadah, S. Chmiela, K.-R. Müller, and A. Tkatchenko, "Quantum-chemical insights from deep tensor neural networks," *Nat. Commun.* **8**, 13890 (2017).
- ²⁶F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, "Bypassing the Kohn-Sham equations with machine learning," *Nat. Commun.* **8**, 872 (2017).
- ²⁷Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. Pande, "MoleculeNet: A benchmark for molecular machine learning," *Chem. Sci.* **9**, 513 (2018).
- ²⁸T. T. Nguyen, E. Székely, G. Imbalzano, J. Behler, G. Csányi, M. Ceriotti, A. W. Götz, and F. Paesani, "Comparison of permutationally invariant polynomials, neural networks, and Gaussian approximation potentials in representing water interactions through many-body expansions," *J. Chem. Phys.* **148**, 241725 (2018).
- ²⁹K. Yao, J. E. Herr, D. W. Toth, R. McKintyre, and J. Parkhill, "The TensorMol-0.1 model chemistry: A neural network augmented with long-range physics," *Chem. Sci.* **9**, 2261–2269 (2018).
- ³⁰S. Fujikake, V. L. Deringer, T. H. Lee, M. Krynski, S. R. Elliott, and G. Csányi, "Gaussian approximation potential modeling of lithium intercalation in carbon nanostructures," *J. Chem. Phys.* **148**, 241714 (2018).
- ³¹H. Li, C. Collins, M. Tanha, G. J. Gordon, and D. J. Yaron, "A density functional tight binding layer for deep learning of chemical Hamiltonians," *J. Chem. Theory Comput.* **14**, 5764–5776 (2018).
- ³²A. Grisafi, A. Fabrizio, B. Meyer, D. M. Wilkins, C. Corminboeuf, and M. Ceriotti, "Transferable machine-learning model of the electron density," *ACS Cent. Sci.* **5**, 57–64 (2019).
- ³³L. Zhang, J. Han, H. Wang, R. Car, and W. E, "Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics," *Phys. Rev. Lett.* **120**, 143001 (2018).
- ³⁴R. T. McGibbon, A. G. Taube, A. G. Donchev, K. Siva, F. Hernández, C. Hargus, K.-H. Law, J. L. Klepeis, and D. E. Shaw, "Improving the accuracy of Møller-Plesset perturbation theory with neural networks," *J. Chem. Phys.* **147**, 161725 (2017).
- ³⁵M. Welborn, L. Cheng, and T. F. Miller III, "Transferability in machine learning for electronic structure via the molecular orbital basis," *J. Chem. Theory Comput.* **14**, 4772–4779 (2018).
- ³⁶G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, "Machine learning of molecular electronic properties in chemical compound space," *New J. Phys.* **15**, 095003 (2013).
- ³⁷L. C. Blum and J.-L. Reymond, "970 million druglike small molecules for virtual screening in the chemical universe database GDB-13," *J. Am. Chem. Soc.* **131**, 8732 (2009).
- ³⁸R. K. Nesbet, "Brueckner's theory and the method of superposition of configurations," *Phys. Rev.* **109**, 1632 (1958).
- ³⁹A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry* (Dover, Mineola, 1996), pp. 231–239.
- ⁴⁰C. Möller and M. S. Plesset, "Note on an approximation treatment for many-electron systems," *Phys. Rev.* **46**, 618 (1934).
- ⁴¹G. Knizia, "Intrinsic atomic orbitals: An unbiased bridge between quantum theory and chemical concepts," *J. Chem. Theory Comput.* **9**, 4834 (2013).
- ⁴²J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.* **98**, 146401 (2007).
- ⁴³S. F. Boys, "Construction of some molecular orbitals to be approximately invariant for changes from one molecule to another," *Rev. Mod. Phys.* **32**, 296–299 (1960).
- ⁴⁴U. Kaldor, "Localized orbitals for NH₃, C₂H₄, and C₂H₂," *J. Chem. Phys.* **46**, 1981–1987 (1967).
- ⁴⁵L. Breiman, "Random forests," *Mach. Learn.* **45**, 5–32 (2001).
- ⁴⁶L. Breiman, "Statistical modeling: The two cultures," *Stat. Sci.* **16**, 199–215 (2001).
- ⁴⁷R. Tripathy, I. Bilonis, and M. Gonzalez, "Gaussian processes with built-in dimensionality reduction: Applications to high-dimensional uncertainty propagation," *J. Comput. Phys.* **321**, 191–223 (2016).
- ⁴⁸See <https://github.com/thomasfmiller/MOB-ML> for the available code.
- ⁴⁹Y. Shao, Z. Gan, E. Epifanovsky, A. T. Gilbert, M. Wormit, J. Kussmann, A. W. Lange, A. Behn, J. Deng, X. Feng, D. Ghosh, M. Goldey, P. R. Horn, L. D. Jacobson, I. Kaliman, R. Z. Khaliullin, T. Kus, A. Landau, J. Liu, E. I. Proynov, Y. M. Rhee, R. M. Richard, M. A. Rohrdanz, R. P. Steele, E. J. Sundstrom, H. L. Woodcock, P. M. Zimmerman, D. Zuev, B. Albrecht, E. Alguire, B. Austin, G. J. O. Beran, Y. A. Bernard, E. Berquist, K. Brandhorst, K. B. Bravaya, S. T. Brown, D. Casanova, C.-M. Chang, Y. Chen, S. H. Chien, K. D. Closser, D. L. Crittenden, M. Diedenhofen, R. A. DiStasio, H. Do, A. D. Dutoi, R. G. Edgar, S. Fatehi, L. Fusti-Molnar, A. Ghysels, A. Golubeva-Zadorozhnaya, J. Gomes, M. W. Hanson-Heine, P. H. Harbach, A. W. Hauser, E. G. Hohenstein, Z. C. Holden, T.-C. Jagau, H. Ji, B. Kaduk, K. Khistyayev, J. Kim, J. Kim, R. A. King, P. Klunzinger, D. Kosenkov, T. Kowalczyk, C. M. Krauter, K. U. Lao, A. D. Laurent, K. V. Lawler, S. V. Levchenko, C. Y. Lin, F. Liu, E. Livshits, R. C. Lochan, A. Luenser, P. Manohar, S. F. Manzer, S.-P. Mao, N. Mardirossian, A. V. Marenich, S. A. Maurer, N. J. Mayhall, E. Neuscamman, C. M. Oana, R. Olivares-Amaya, D. P. O'Neill, J. A. Parkhill, T. M. Perrine, R. Peverati, A. Prociuk, D. R. Rehn, E. Rosta, N. J. Russ, S. M. Sharada, S. Sharma, D. W. Small, A. Sodt, T. Stein, D. Stück, Y.-C. Su, A. J. Thom, T. Tsuchimochi, V. Vanovschi, L. Vogt, O. Vydrov, T. Wang, M. A. Watson, J. Wenzel, A. White, C. F. Williams, J. Yang, S. Yeganeh, S. R. Yost, Z.-Q. You, I. Y. Zhang, X. Zhang, Y. Zhao, B. R. Brooks, G. K. Chan, D. M. Chipman, C. J. Cramer, W. A. Goddard, M. S. Gordon, W. J. Hehre, A. Klamt, H. F. Schaefer, M. W. Schmidt, C. D. Sherrill, D. G. Truhlar, A. Warshel, X. Xu, A. Aspuru-Guzik, R. Baer, A. T. Bell, N. A. Besley, J.-D. Chai, A. Dreuw, B. D. Dunietz, T. R. Furlani, S. R. Gwaltney, C.-P. Hsu, Y. Jung, J. Kong, D. S. Lambrecht, W. Liang, C. Ochsenfeld, V. A. Rassolov, L. V. Slipchenko, J. E. Subotnik, T. Van Voorhis, J. M. Herbert, A. I. Krylov, P. M. Gill, and M. Head-Gordon, "Advances in molecular quantum chemistry contained in the Q-Chem 4 program package," *Mol. Phys.* **113**, 184 (2015).
- ⁵⁰S. H. Vosko, L. Wilk, and M. Nusair, "Accurate spin-dependent electron liquid correlation energies for local spin density calculations: A critical analysis," *Can. J. Phys.* **58**, 1200 (1980).
- ⁵¹C. Lee, W. Yang, and R. G. Parr, "Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density," *Phys. Rev. B* **37**, 785 (1988).
- ⁵²A. D. Becke, "Density-functional thermochemistry. III. The role of exact exchange," *J. Chem. Phys.* **98**, 5648 (1993).
- ⁵³P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, "Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields," *J. Phys. Chem.* **98**, 11623 (1994).
- ⁵⁴P. C. Hariharan and J. A. Pople, "The influence of polarization functions on molecular orbital hydrogenation energies," *Theor. Chim. Acta* **28**, 213 (1973).
- ⁵⁵G. Bussi and M. Parrinello, "Accurate sampling using Langevin dynamics," *Phys. Rev. E* **75**, 056707 (2007).
- ⁵⁶H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, M. Schütz, P. Celani, W. Györfy, D. Kats, T. Korona, R. Lindh, A. Mitrushenkov, G. Rauhut, K. R. Shamasundar, T. B. Adler, R. D. Amos, S. J. Bennie, A. Bernhardsson, A. Berning, D. L. Cooper, M. J. O. Deegan, A. J. Dobbyn, F. Eckert, E. Goll, C. Hampel, A. Hesselmann, G. Hetzer, T. Hrenar, G. Jansen, C. Köppl, S. J. R. Lee, Y. Liu, A. W. Lloyd, Q. Ma, R. A. Mata, A. J. May, S. J. McNicholas, W. Meyer, T. F. Miller III, M. E. Mura, A. Nicklass, D. P. O'Neill, P. Palmieri, D. Peng, K. Pflüger, R. Pitzer, M. Reiher, T. Shiozaki, H. Stoll, A. J. Stone, R. Tarroni, T. Thorsteinsson, M. Wang, and M. Welborn, *MOLPRO*, version 2018.3, a package of *ab initio* programs, 2018, see <http://www.molpro.net>.
- ⁵⁷T. H. Dunning, "Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen," *J. Chem. Phys.* **90**, 1007 (1989).
- ⁵⁸S. Saebø and P. Pulay, "Local treatment of electron correlation," *Annu. Rev. Phys. Chem.* **44**, 213–236 (1993).
- ⁵⁹J. Čížek, "On the correlation problem in atomic and molecular systems. Calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods," *J. Chem. Phys.* **45**, 4256 (1966).

- ⁶⁰C. Hampel and H. J. Werner, "Local treatment of electron correlation in coupled cluster theory," *J. Chem. Phys.* **104**, 6286–6297 (1996).
- ⁶¹R. J. Bartlett, J. D. Watts, S. A. Kucharski, and J. Noga, "Non-iterative fifth-order triple and quadruple excitation energy corrections in correlated methods," *Chem. Phys. Lett.* **165**, 513–522 (1990).
- ⁶²M. Schütz, "Low-order scaling local electron correlation methods. III. Linear scaling local perturbative triples correction (T)," *J. Chem. Phys.* **113**, 9986–10001 (2000).
- ⁶³R. Polly, H.-J. Werner, F. R. Manby, and P. J. Knowles, "Fast Hartree-Fock theory using local density fitting approximations," *Mol. Phys.* **102**, 2311–2321 (2004).
- ⁶⁴C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
- ⁶⁵GPpy, GPpy: A gaussian process framework in python, <http://github.com/SheffieldML/GPpy>, since 2012.
- ⁶⁶In principle, the smoothness of the Matérn kernel could be taken as a kernel hyperparameter; however, this possibility was not explored in this work.
- ⁶⁷F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.* **12**, 2825 (2011), <http://www.jmlr.org/papers/v12/pedregosa11a.html>.
- ⁶⁸C. Cortes, L. D. Jackel, S. A. Solla, V. Vapnik, and J. S. Denker, "Learning curves: Asymptotic values and rate of convergence," in *Advances in Neural Information Processing Systems 6*, edited by J. D. Cowan, G. Tesauro and J. Alspector (Morgan-Kaufmann, 1994), pp. 327–334.
- ⁶⁹K. Pearson, "Mathematical contributions to the theory of evolution. III. Regression, heredity, and panmixia," *Philos. Trans. R. Soc., A* **187**, 253–318 (1896).
- ⁷⁰F. A. Faber, A. S. Christensen, B. Huang, and O. A. von Lilienfeld, "Alchemical and structural distribution based representation for universal quantum machine learning," *J. Chem. Phys.* **148**, 241717 (2018).
- ⁷¹A. S. Christensen, F. A. Faber, B. Huang, L. A. Bratholm, A. Tkatchenko, K. R. Muller, O. A. von Lilienfeld, "QML: A Python toolkit for quantum machine learning," <https://github.com/qmlcode/qml> (2017).
- ⁷²B. Kramer and A. MacKinnon, "Localization: Theory and experiment," *Rep. Prog. Phys.* **56**, 1469 (1993).
- ⁷³L. Cheng, M. Welborn, A. S. Christensen, and T. F. Miller III, "Thermalized (350 K) QM7b, GDB-13, water, and short alkane quantum chemistry dataset including MOB-ML features," CaltechDATA dataset, <https://doi.org/10.22002/d1.1177> (2019).