

# Alchemical and structural distribution based representation for universal quantum machine learning

Cite as: J. Chem. Phys. **148**, 241717 (2018); <https://doi.org/10.1063/1.5020710>

Submitted: 26 December 2017 • Accepted: 28 February 2018 • Published Online: 20 March 2018

Felix A. Faber,  Anders S. Christensen, Bing Huang, et al.

## COLLECTIONS

Paper published as part of the special topic on [Data-Enabled Theoretical Chemistry](#)



View Online



Export Citation



CrossMark

## ARTICLES YOU MAY BE INTERESTED IN

[Atom-density representations for machine learning](#)

The Journal of Chemical Physics **150**, 154110 (2019); <https://doi.org/10.1063/1.5090481>

[Less is more: Sampling chemical space with active learning](#)

The Journal of Chemical Physics **148**, 241733 (2018); <https://doi.org/10.1063/1.5023802>

[Machine learning for interatomic potential models](#)

The Journal of Chemical Physics **152**, 050902 (2020); <https://doi.org/10.1063/1.5126336>

Lock-in Amplifiers  
up to 600 MHz



Zurich  
Instruments



# Alchemical and structural distribution based representation for universal quantum machine learning

Felix A. Faber, Anders S. Christensen, Bing Huang, and O. Anatole von Lilienfeld<sup>a)</sup>

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

(Received 26 December 2017; accepted 28 February 2018; published online 20 March 2018)

We introduce a representation of any atom in any chemical environment for the automatized generation of universal kernel ridge regression-based quantum machine learning (QML) models of electronic properties, trained throughout chemical compound space. The representation is based on Gaussian distribution functions, scaled by power laws and explicitly accounting for structural as well as elemental degrees of freedom. The elemental components help us to lower the QML model's learning curve, and, through interpolation across the periodic table, even enable "alchemical extrapolation" to covalent bonding between elements not part of training. This point is demonstrated for the prediction of covalent binding in single, double, and triple bonds among main-group elements as well as for atomization energies in organic molecules. We present numerical evidence that resulting QML energy models, after training on a few thousand random training instances, reach chemical accuracy for out-of-sample compounds. Compound datasets studied include thousands of structurally and compositionally diverse organic molecules, non-covalently bonded protein side-chains, (H<sub>2</sub>O)<sub>40</sub>-clusters, and crystalline solids. Learning curves for QML models also indicate competitive predictive power for various other electronic ground state properties of organic molecules, calculated with hybrid density functional theory, including polarizability, heat-capacity, HOMO-LUMO eigenvalues and gap, zero point vibrational energy, dipole moment, and highest vibrational fundamental frequency. © 2018 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/1.5020710>

## I. INTRODUCTION

Ground-state properties of chemical compounds can generally be estimated with acceptable accuracy using methods such as *ab initio* quantum chemistry or density functional theory (DFT).<sup>1</sup> However, these can be computationally expensive and therefore have a limited applicability, especially for larger systems. Alternatively, inductive machine learning of quantum mechanical properties, i.e., quantum machine learning (QML) models, can infer properties directly or even predict the electron density which in turn can be used to calculate properties,<sup>2</sup> by training on large datasets of reference property/compound pairs. QML models can have an exceptional trade-off between predictive accuracy and computational cost. For example, in 2017 we showed that QML models can estimate hybrid DFT atomization energies as well as several other properties of medium-sized organic molecules with prediction errors lower than chemical accuracy (~0.04 eV)—multiple orders of magnitude faster than hybrid DFT.<sup>3</sup>

The system variable defining the ground-state properties of a given compound is its external potential, a simple function of interatomic distances and nuclear charges. However, using this information directly to measure similarity results

in QML models with rather disappointing predictive power. This can be mitigated by the transformation of system variables into "representations." Such transformations can either be designed by human intuition or be included in the learning problem, e.g., when using neural networks (NNs) which include representation learning in the supervised learning task. Letting a NN find the representation has proven to yield models with low out-of-sample prediction errors.<sup>4–6</sup> This approach, however, has the drawback that the representation and model are intermingled within the NN, making it less amenable to human understanding, interpretation, adaptation, and further improvement. Furthermore, such machine designed representations do not necessarily lead to better QML performance than representations designed by humans (*vide infra*).

There are many ways of manually encoding the 3D structure and chemical composition of a compound into a suitable representation. For example, we can represent a compound as a list of interatomic potentials.<sup>7–9</sup> Another approach consists of creating a fingerprint of the compound, transforming internal coordinates into a fixed set of numbers. For example, this can be done by projecting the coordinates on to a set of basis functions<sup>10</sup> or by using the topology of the structure.<sup>11</sup> Distributions of internal coordinates represent another systematic approach, shown to yield well performing QML models applicable throughout chemical space.<sup>12,13</sup> Additional use of bags containing angular and dihedral distributions has led to further improvements in resulting QML models.<sup>3,9,14,15</sup> Bagging

<sup>a)</sup>anatole.vonlilienfeld@unibas.ch



based on atom types, however, severely hinders resulting QML models from transferring what has been learned from one atom type to another—a desirable feature for chemically diverse systems.

In this work, we introduce a new atomic environment representation, with two key differences to previous distribution-based work. (i) The representation is not binned by atomic species. Instead, compositional information is encoded directly into the distributions. This allows measuring not only structural differences but also “alchemical” differences between elements in the atomic environments. The idea of computational alchemy, amounting to continuous interpolation of Hamiltonians of two different systems, is well established in quantum chemistry and statistical mechanics and can be exploited for virtual exploration campaigns in chemical space with increased efficiency.<sup>16</sup> Recently, it has been shown that alchemical estimates of covalent bond potentials can even surpass generalized gradient approximated DFT accuracy.<sup>17</sup> The foundation of a continuous chemical space has been reviewed previously.<sup>12</sup> Alchemical distance measures in the context of QML were already exploited previously when using the Coulomb matrix (CM),<sup>7</sup> the Fourier series distribution based representations,<sup>18</sup> the Faber, Lindmaa, Lilienfeld, Armiento (FLLA) crystal representation,<sup>19</sup> and within smooth overlap of atomic potential (SOAP) representations.<sup>20</sup> For this work, we have identified a new functional form with improved performance due to alchemical contributions to the distance measure. (ii) We use a set of multidimensional distributions of interatomic many-body expansions rather than several 1D bins of internal coordinates. The distributions are built recursively so that an  $m$ -body distribution contains the same information as the  $(m - 1)$ -body distribution plus additional  $m$ -body information. This particular combination combines similarity to the potential energy target function and compliance with many known (translational, rotational, permutational) invariances.

## II. THEORY

In this section, we first motivate the ideas that have led to this study. Thereafter, we discuss the functional form and the variational degrees of freedom, which we have introduced, as well as the resulting compound distances. Then, an analysis of the functional form is performed using the molecule water as an example. Finally, numerical results for parameters’ optimization runs are discussed.

### A. Kernel ridge regression

In order to profit from robustness, ease of error convergence, computational efficiency, and simplicity, we base our studies preferably on kernel ridge regression (KRR) models.<sup>21–24</sup> However, we consider this rather a question of taste and believe that other regressors, such as neural networks, will produce similar results if properly converged.

KRR estimates property  $p$  of query compound  $\mathbf{C}$  as a weighted sum of kernel basis functions placed on each of  $N$  training compounds  $\{\mathbf{C}_k\}$ ,

$$p^{\text{est}}(\mathbf{C}) = \sum_{k=1}^N \alpha_k K(\mathbf{C}, \mathbf{C}_k), \quad (1)$$

$$\boldsymbol{\alpha} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{p}^{\text{train}}, \quad (2)$$

where the solution for the weights  $\{\alpha_k\}$  is obtained through linear regression with regularizer  $\lambda$  (typically negligibly small because of the absence of noise in training data obtained from quantum chemistry calculations).

Throughout this work we rely on atomistic (scalable) Gaussian kernels,  $K(\mathbf{C}, \mathbf{C}') = \sum_{I \in \mathbf{C}} \sum_{J \in \mathbf{C}'} k(\Delta(\mathcal{A}_M(I), \mathcal{A}_M(J)))$ , as already used in Refs. 14 and 25–27. As such, KRR renders the selection of a functional form which represents an atom in its chemical environment mandatory. Obviously, this choice is fundamentally related to our understanding of chemistry and is known to dramatically affect the performance of resulting QML models, see, e.g., Refs. 3 and 9. It is for this reason that we draw our inspiration from the fundamental laws of quantum mechanics which specify the definition of system (Hamiltonian) and property (observable) and which spell out the numerical recipe which links the two.<sup>1</sup>

The genesis of this study is due to the fact that the total potential energy, the expectation value of a compound’s electronic Hamiltonian, constitutes the central figure of merit for convergence towards the wavefunction by virtue of the variational principle. When considering Eq. (2), it should be obvious that the kernel (and thereby representation) is independent of a specific property. Units and property dependence are introduced through the regression weights only. This has also already been demonstrated numerically for multiple properties using the same kernel.<sup>28</sup> As such, the role of the kernel is reminiscent of the wavefunction which can be used to predict arbitrarily many observables by evaluating the expectation values of the corresponding operators, always using the same wavefunction: Once the kernel is inverted, arbitrarily many sets of regression coefficients can easily be generated provided that their corresponding property reference values are known. The Hamiltonian’s expectation value, i.e., the potential energy, therefore governs the shape of the wavefunction. We therefore assume that a representation, optimized for energy predictions only, is fundamentally more advantageous than representations obtained by minimizing prediction errors of some integrated observable. Consequently, the focus employed in this study has been to identify a representation which is inspired by the energy changes occurring due to changes in chemical composition and covalent and non-covalent bonding. The accuracy of quantum mechanics when predicting other properties (observables) as expectation values of operators depends crucially on the quality of the wavefunction. Here, we follow a similar argument: The better the representation the better the energy prediction, implying that energy prediction errors should be minimized in the functional space of the representation.

### B. Representation

We use a set of interatomic  $M$ -body expansions  $\mathcal{A}_M(I) = \{A_1(I), A_2(I), A_3(I), \dots, A_M(I)\}$  which contain up to  $M$ -body interactions to represent the structural and chemical

environment of an atom  $I$  in compound  $C$ .  $A_m(I)$  is a weighted sum that runs over all  $m$ -body interactions. Each element in the sums consists of Gaussian basis functions, placed on structural and elemental degrees of freedom and multiplied by a scaling function  $\xi_m$ . Structural values encode geometrical information about the system, such as interatomic distances or angles. As elemental parameters, we use the period  $P$  and group  $G$  from the periodic table. The scaling functions  $\xi_m$  are used to weigh the importance of each Gaussian, based on internal system coordinates. We now consider only the first three distributions in  $\mathcal{A}_M(I)$  for an atom  $I$ . We have also derived, implemented, and tested the 4-body  $A_4(I)$  distributions. However, the predictive accuracy improvements of resulting QML models were found to be negligible in comparison to the 3-body expansion. Also, the computational cost for generating large kernel matrices increases substantially when going from third to fourth order terms. We refer to the [supplementary material](#) for the derivation.

The first-order expansion  $A_1(I)$  accounts for chemical composition (stoichiometry) and is modeled by a Gaussian function placed on period  $P_I$  and group  $G_I$  in the periodic table of element  $I$ ,

$$A_1(I) = \mathcal{N}(\mathbf{x}_I^{(1)}) = e^{-\frac{(P_I - \chi_1)^2}{2\sigma_P^2} - \frac{(G_I - \chi_2)^2}{2\sigma_G^2}}, \quad (3)$$

where  $\mathbf{x}_I^{(1)} = \{P_I, \sigma_P; G_I, \sigma_G\}$ , with respective widths  $\sigma_P$  and  $\sigma_G$ .  $\sigma_P$  and  $\sigma_G$  can be seen as elemental smearing parameters, which control the near-sightedness of elements in the periodic table.  $\chi_1$  and  $\chi_2$  represent dummy variables for the period and group, to be integrated out when evaluating the Euclidean distance [see Eq. (4)]. For  $A_1(I)$ , the scaling function is set to unity since stoichiometry is geometry independent. We are not aware of other representations in the literature which employ similar distribution functions in the periodic table.

$A_2(I)$  is a product of  $A_1(I)$  and a sum that runs over all neighboring atoms  $i$ :  $A_2(I) = \mathcal{N}(\mathbf{x}_I^{(1)}) \sum_{i \neq I} \mathcal{N}(\mathbf{x}_{iI}^{(2)}) \xi_2(d_{iI})$ ,  $\mathbf{x}_{iI}^{(2)} = \{d_{iI}, \sigma_d; P_i, \sigma_P; G_i, \sigma_G\}$ , where  $d_{iI}$  and  $\sigma_d$  correspond

to the interatomic distance at which a Gaussian is placed and its width, respectively.  $\xi_2$  corresponds to the 2-body, interatomic distance dependent, scaling function which takes the form of the power laws discussed below. Note that letting  $\sigma_P$  and  $\sigma_G$  approach zero is equivalent to using a radial distribution function (RDF) for each element pair. This attribute of the representation holds for any of  $A_m(I)$ , i.e.,  $\sigma_P, \sigma_G \rightarrow 0$  is equivalent to creating a separate distribution for each chemical element  $m$ -tuple in  $A_m(I)$ .  $A_3(I)$  is the logical extension from  $A_2(I)$ ; it has a different scaling function with an additional summation, running over all neighboring atoms  $j$ :  $A_3(I) = \mathcal{N}(\mathbf{x}_I^{(1)}) \sum_{i \neq I} \mathcal{N}(\mathbf{x}_{iI}^{(2)}) \sum_{j \neq i, I} \mathcal{N}(\mathbf{x}_{ijI}^{(3)}) \xi_3(d_{iI}, d_{jI}, \theta_{ij}^I)$ ,  $\mathbf{x}_{ijI}^{(3)} = \{\theta_{ij}^I, \sigma_\theta; P_j, \sigma_P; G_j, \sigma_G\}$ .  $P_j$  and  $G_j$ , similar to  $P_i$  and  $G_i$ , correspond to the period and group of atom  $j$ . Again,  $\xi_3(d_{iI}, d_{jI}, \theta_{ij}^I)$  is the (three-body) scaling function, and  $\theta_{ij}^I$  is the principal angle between the two distance vectors  $\vec{r}_{iI}$  and  $\vec{r}_{jI}$  which span from  $I$  to  $i$  and  $I$  to  $j$ , respectively.  $\sigma_\theta$  is the width of the Gaussian placed at  $\theta_{ij}^I$ . Letting  $\sigma_d$  go to infinity in  $A_3$  is equivalent to using a type of angular distribution function (ADF), which in one form or another has already been used in several representations.<sup>3,14,15</sup>  $A_3$  can therefore be seen as a generalized ADF containing more structural information. Figure 1 illustrates how  $A_3(I)$  looks for hydrogen, carbon, and oxygen atoms in ethanol.

The scaling functions  $\xi$  we have chosen for this work correspond to simple power laws. They have been modified from the leading order two- and three-body dispersion laws by London,  $1/r^6$ , and Axilrod-Teller-Muto,<sup>29,30</sup>  $1/r^9$ . Such dispersion expressions were previously already used by some of us.<sup>14</sup> Our scaling functions, however, use different exponents for the radial decay and set the  $C_6$  and  $C_9$  coefficients to unity, as early tests indicated better performance for this choice. For periodic systems, however, a very large cutoff radius would be needed in order to converge the distances between two atomic environments, when using the optimized exponents. We have therefore augmented the scaling functions by a previously used soft cutoff function,<sup>31</sup> which goes to zero at 9 Å.




FIG. 1. The three-body term ( $A_3$ ) as a function of radial ( $d$ ) and angular ( $\theta$ ) degrees of freedom in the atomic environments of O, C, and H (circled) in ethanol. For simplicity, we show the three-body term without elemental smearing where it reduces to a number of two-dimensional distributions for each element triple.

### C. Distances and scalar products

In order to train and evaluate the KRR model in Eq. (1), proper distance measures must be specified. We have found good performance when using a weighted sum of the distances between each  $m$ -body expansion  $\Delta(\mathcal{A}_M(I), \mathcal{A}_M(J))^2 \equiv \sum_{m=0}^M \beta_m \Delta(A_m(I), A_m(J))^2$  as a distance between two atomic environments  $\mathcal{A}_M(I)$  and  $\mathcal{A}_M(J)$ . Here,  $\beta_m$  is another hyperparameter which weighs the importance of each expansion order.

$$\Delta(A_m(I), A_m(J))^2 = \frac{1}{\zeta_m^2} \int_{\mathbb{R}^{3m-1}} d\chi_1 \cdots d\chi_{3m-1} (A_m(I) - A_m(J))^2, \quad (4)$$

$$\frac{1}{\zeta_1^2} \int_{\mathbb{R}^2} d\chi_1 d\chi_2 A_1(I) A_1(J) = \frac{1}{2} \exp\left(-\frac{(P_I - P_J)^2}{4\sigma_P^2} - \frac{(G_I - G_J)^2}{4\sigma_G^2}\right), \quad (5)$$

$$\begin{aligned} \frac{1}{\zeta_2^2} \int_{\mathbb{R}^5} d\chi_1 \cdots d\chi_5 A_2(I) A_2(J) &= \frac{1}{2\sqrt{2}} \exp\left(-\frac{(P_I - P_J)^2}{4\sigma_P^2} - \frac{(G_I - G_J)^2}{4\sigma_G^2}\right) \\ &\times \sum_{i \neq l}^{n_I} \xi_2(d_{il}) \sum_{j \neq J}^{n_J} \exp\left(-\frac{(d_{jJ} - d_{il})^2}{4\sigma_d^2} - \frac{(P_i - P_j)^2}{4\sigma_P^2} - \frac{(G_i - G_j)^2}{4\sigma_G^2}\right) \xi_2(d_{jJ}), \end{aligned} \quad (6)$$

$$\begin{aligned} \frac{1}{\zeta_3^2} \int_{\mathbb{R}^8} d\chi_1 \cdots d\chi_8 A_3(I) A_3(J) &= \frac{1}{16} \exp\left(-\frac{(P_I - P_J)^2}{4\sigma_P^2} - \frac{(G_I - G_J)^2}{4\sigma_G^2}\right) \\ &\times \sum_{i \neq l}^{n_I} \sum_{j \neq J}^{n_J} \exp\left(-\frac{(d_{jJ} - d_{il})^2}{4\sigma_d^2} - \frac{(P_i - P_j)^2}{4\sigma_P^2} - \frac{(G_i - G_j)^2}{4\sigma_G^2}\right) \\ &\times \sum_{k \neq i, l}^{n_I} \xi_2(d_{il}, d_{kl}, \theta_{ik}^I) \sum_{l \neq j, J}^{n_J} \exp\left(-\frac{(\theta_{ik}^I - \theta_{jl}^J)^2}{4\sigma_\theta^2} - \frac{(P_k - P_l)^2}{4\sigma_P^2} - \frac{(G_k - G_l)^2}{4\sigma_G^2}\right) \xi_3(d_{jJ}, d_{il}, \theta_{jk}^J). \end{aligned} \quad (7)$$

Note that third and fourth order terms become prohibitively expensive to calculate directly. However, this can be circumvented to a large extent by slightly modifying the distributions and solving the angular integrals in Fourier space. Further details about the corresponding equations and derivations can also be found in the [supplementary material](#).

### D. Comparison to other distribution based representations

Probably the largest difference in how  $\mathcal{A}$  represents nuclear configurations, when compared to many of the previously published distribution based representations, lies in the 3-body term (since  $A_2$  is a radial distribution function if  $\sigma_P$  and  $\sigma_G$  go to zero). In this subsection, we highlight the differences between  $A_3$ , or conventional ADF or RDF for representing the structure of the water molecule.

As ADF, we use  $A_3$  with the limit  $\sigma_d \rightarrow \infty$  and we model RDF by  $A_2$ . Furthermore, no scaling function ( $\xi_2 = \xi_3 = 1$ ) is used and we let  $\sigma_P$  and  $\sigma_G$  go to zero since we only examine how representations distinguish structural differences among different geometries of the water molecule. This results in  $A_3$  and ADF being  $\sum_{i \neq l} \mathcal{N}(d_{il}, \sigma_d) \sum_{j \neq i, l} \mathcal{N}(\theta_{ij}^I, \sigma_\theta)$  and  $\sum_{i \neq l} \sum_{j \neq i, l} \mathcal{N}(\theta_{ij}^I, \sigma_\theta)$ , respectively, for each element triple and RDF being  $\sum_{i \neq l} \mathcal{N}(d_{il}, \sigma_d)$  for each element pair.

The distances between each distribution term are evaluated as Euclidean ( $L_2$ ) norms, as shown in Eq. (4).  $\zeta_m$  are normalization constants, which ensure that all individual basis functions integrate to 1 in the  $L_2$ -norm. All integrals can be solved analytically since they consist of a sum of Gaussian products. The explicit form of the  $A_m$  integrals for  $m = 1, \dots, 3$  is shown in Eqs. (5)–(7). Details on how to evaluate the  $A_3$  and  $A_4$  integrals in Fourier space can be found in the [supplementary material](#),

Figure 2 shows how the distance measure between two water molecules changes as one distorts the geometry of one of them away from its equilibrium structure. Both, RDF and ADF result for oxygen as well as for H in large configurational domains with substantially zero distance to the minimum, implying a severe lack of sensitivity.  $A_3$  by contrast produces a qualitatively meaningful picture with a single well-defined well around the minimum.

We have also studied the performance for modeling the energy of the water molecule. In Fig. 3, the training error for atomization energies is shown for a linear kernel KRR model with  $A_3$ , ADF, RDF, or RDF + ADF as representations. The linear kernel is used as a difficult test in how far representations can model a nonlinear property, such as the energy, in terms of linear basis functions. The errors are significantly lower when using  $A_3$  instead of the other representations, including RDF + ADF. Generally, potential energy surfaces of a three-atom system cannot be decomposed into a sum of functions of only one internal coordinate (internuclear distance  $\mathbf{d}$  or angle  $\theta$ ). That is,  $E(\mathbf{d}, \theta) \neq E(\mathbf{d}) + E(\theta)$ . Using an ADF, an RDF, or a linear combination of the two however would result precisely in such a model, as well as most force fields. This also explains the relatively large errors for these representations as well as the unreliable performance of pair-wise potentials when it comes to distorted molecules.




FIG. 2. Heat maps of normalized  $L_2$  distances, using three different (not yet scaled) representations (RDF, ADF, and our new representation). The color code from black to white indicates normalized distance, ranging from 0 to 1, respectively. The distances are measured between a reference water molecule, in equilibrium geometry (cross), and a distorted water molecule. Distances are measured separately for the oxygen (left) and hydrogen atoms (right). The distorted water molecule is generated by uniformly stretching both OH bonds ( $d_{\text{OH1}} = d_{\text{OH2}} = l$ ) and bending the HOH angle ( $\phi$ ) of the reference molecule. The relevant equations for the three representations are given in Sec. II D.

$A_3$  on the other hand does not decouple distances and angles and can, by construction, model any three-body potential.

These observations give insight as to why our new representation performs better than the other distribution based

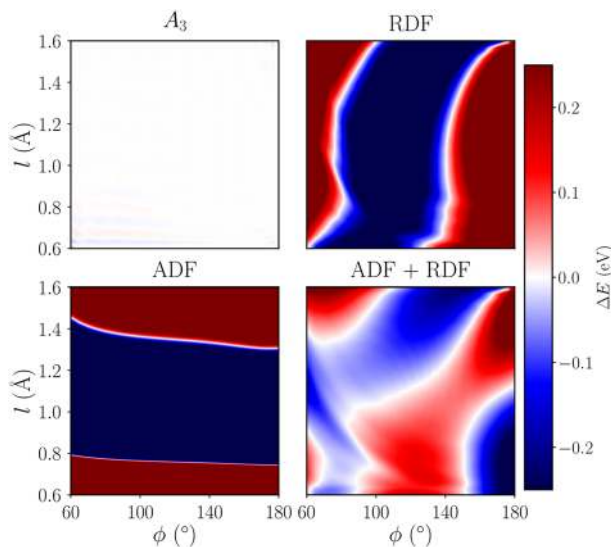


FIG. 3. Heat maps depicting the signed atomization energy error of a water molecule using the same coordinate system as in Fig. 2. The errors correspond to linear kernels in KRR fitted to DFT calculated energies (PBE/def2svp). Four representations have been used: Top left: our new  $A_3$ . Top right: radial distribution function (RDF) for each element pair. Bottom left: angular distribution function (ADF) for each element triple. Bottom right: RDF + ADF. The training data consist of an equidistant grid of  $50 \times 50$  points along  $l$  and  $\phi$  within the range of the figures.

representations: Using ADF's and RDF's as representations one might be able to capture slices of the many-body picture, the fact that there is a linear mapping between  $A_n$  and a  $n$ -body potential energy surface, however, appears to make it easier to improve the performance also for non-linear kernels.

## E. Optimization

### 1. Hyperparameters

The use of our representation in combination with KRR yields multiple hyperparameters. While one could, in principle, attempt to optimize all of them, using several datasets, and efficient optimizers, such as gradient, Monte Carlo, genetic, or simplex methods, we have found that the problem is sensitive only to a small subset of parameters. As such, the exact choice of many hyperparameters is not critical for the out-of-sample errors, and resulting models perform typically well as long as values are used which have a similar order of magnitude. Unless otherwise specified, the following hyperparameter values have been used:  $\sigma_P = \sigma_G = 1.6$ ,  $\sigma_d = 0.2 \text{ \AA}$ ,  $\sigma_\theta = \pi$  black,  $\beta_1 = 2$ ,  $\beta_2 = \sqrt{8}$ ,  $\beta_3 = 1.6$ . For the water cluster and the protein-side chain-side chain interaction (SSI) dataset there is no to little variation in chemical composition, and no elemental smearing has been used.

### 2. Scaling power law parameters

We have screened radial exponents  $n_2$  and  $n_3$  for the scaling functions  $\xi_2(d_{ij}) = \frac{1}{d_{ij}^{n_2}}$  and  $\xi_3(d_{ij}, d_{jl}, \theta_{ij}^l)$

$$= \frac{1 - 3 \cos(\theta_{ij}^l) \cos(\theta_{ij}^i) \cos(\theta_{ij}^j)}{(d_{ij} d_{jl} d_{ij})^{n_3}},$$

using atomization energies for a subset of the QM9 dataset in order to identify the optimal exponents. Corresponding learning curves are shown in Fig. 4. First, we have screened  $\xi_2$ , using  $A_2$  as a representation, yielding the lowest offset for  $n_2 = 4$ . Keeping this exponent for  $\xi_2$  fixed, we then proceeded to screen the exponent  $\xi_3$  in  $A_3$ . We found that  $n_3 = 2$  corresponded to the best exponents for  $\xi_3$ . We have used these values throughout this work, and unless something else is specified, the optimal scaling functions read




FIG. 4. Optimization of exponents in scaling power laws. Left: Out-of-sample MAE for atomization/formation energy predictions as a function of training set size on the QM9 dataset. Learning curves are generated using KRR with  $A_2$  as a representation. The legends indicate the exponent  $n_2$  used in the scaling power law,  $\xi_2(d)$ . Right: Out-of-sample MAE for atomization/formation energy predictions as a function of training set size on the QM9 dataset. Learning curves are generated using KRR with  $A_3$  as a representation. The legends indicate the exponent  $n_3$  used in the scaling power law,  $\xi_3(d)$ .

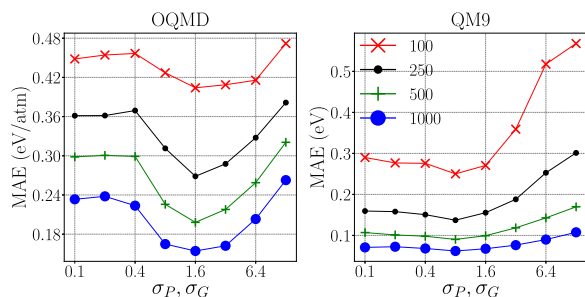


FIG. 5. Changes in out-of-sample MAE as a function of uniform Gaussian width ( $\sigma_P$  and  $\sigma_G$ ) used for elemental smearing. Results for energy predictions in the OQMD (left) and QM9 (right) datasets, respectively. Legends indicate the training set size.

$$\xi_2(d_{il}) = \frac{1}{d_{il}^4},$$

$$\xi_3(d_{il}, d_{jl}, \theta_{ij}^l) = \frac{1 - 3 \cos(\theta_{ij}^l) \cos(\theta_{ij}^i) \cos(\theta_{il}^j)}{(d_{il}d_{jl}d_{ij})^2}. \quad (8)$$

### 3. Alchemical smearing

Parameters associated with the elemental smearing have a strong effect on the predictive power of the QML models. We have screened the corresponding values of  $\sigma_P$  and  $\sigma_G$  using energy prediction errors for the open quantum materials database (OQMD) and QM9 dataset for different training set sizes. These two datasets have been used due to their (relatively) high (OQMD) and low (QM9) chemical diversity in terms of number of differing elements in the dataset. The optimal alchemical Gaussian widths vary only slightly across the two sets, as shown in Fig. 5. A circular Gaussian with width  $\sigma_P = \sigma_G = \sim 1.6$ , which amounts to  $\sim 90\%$  overlap between neighboring elements, corresponds to a relatively deep well with minimal mean absolute error (MAE) for the OQMD dataset, no matter the training set size. The fact that the optimal width stays constant with respect to training set size is beneficial: the elemental smearing can be optimized using relatively small training sets and can then be applied to larger training sets. Comparing the MAE from a model with  $\sigma_P = \sigma_G = 0.1$  (which in practice is equivalent to zero overlap between different atomic types), using the optimal  $\sigma_P = \sigma_G$  lowers the MAE by  $\sim 9.9\%$  for the OQMD dataset at 100 training samples, which increases up to  $\sim 34\%$  when 1k training samples are used. Prediction errors for the QM9 dataset indicate similar behavior, yet much less pronounced. For the largest training set (1000 molecules), the optimization well becomes very shallow, consistent with the lack of compositional diversity in QM9.

Unsurprisingly, datasets with higher chemical diversity benefit more from using the optimized elemental widths. It may therefore not always be beneficial to include any elemental overlap, especially for datasets with low elemental diversity, as it is computationally more expensive to do so.

## III. DATASETS

We have used multiple datasets to benchmark out-of-sample accuracy of energy predictions of our model. These

datasets include organic molecules, crystals, biomolecular dimers, water clusters, and main-group diatomics. Some of the datasets are high-quality, have already been published, and are in widespread use. Additional low quality datasets have been generated, merely in order to accumulate additional evidence for the relative improvement of the new representation. Since test set predictions are always close to zero by construction, we exclusively report prediction errors as out-of-sample errors (averaged through cross-validation) with respect to reference validation numbers. All errors reported correspond to at least 10 cross-validation runs for each training set size.

The reader should note that we only report errors of ML models trained on individual datasets. Simultaneous training on several datasets would introduce significant noise (due to the datasets originating from different levels of theory) and thereby hamper an unequivocal comparative analysis of the results. For future applications and models generated within multi-fidelity ML frameworks,<sup>32</sup> pooling the various datasets to train a single ML model might be more desirable.

### A. Organic molecules: QM9

The QM9 dataset<sup>33</sup> corresponds to hybrid DFT<sup>34</sup> based structures and properties of 134k organic molecules with up to nine atoms (C, O, N, or F), not counting hydrogen. SMILES strings of these molecules correspond to a subset of the GDB-17 dataset.<sup>35</sup> The 3k organic molecules, which fail SMILES consistency tests,<sup>33</sup> were removed before use, i.e., structures where the SMILES strings of the relaxed structure differ from the original GDB-17 SMILES strings.

A random subset of 22k molecules was selected from QM9 for training and testing. 2k molecules were used for testing and up to 20k molecules were used for training.

The datasets were sampled differently from QM9 in Sec. IV B when we investigated how excluding elements from the training set affected the out-of-sample predictions. In total, two tests sets were used, each associated with two training sets. The first test set consisted of molecules containing nitrogen, and the second test set consisted of molecules containing oxygen. The two training sets associated with the first test set consisted of molecules containing nitrogen or not, respectively. The two training sets associated with the second test set consisted of molecules containing oxygen or not, respectively.

### B. Organic molecules: QM7b

Due to widespread use, we also included the more established QM7b dataset.<sup>36</sup> QM7b was also derived from GDB.<sup>37</sup> It contains hybrid DFT (PBE0<sup>38,39</sup>) structures and properties of  $\sim 7k$  organic molecules with up to seven atoms (C, O, N, S, or Cl), not counting H. We have drawn at random up to 5k molecules for training and 2k molecules for testing.

### C. Biomolecular dimers: SSI

For intra-molecular and non-equilibrium interactions, we used a subset of 2356 neutral dimers from the recently published SSI dataset.<sup>40</sup> The SSI dataset is a collection of dimers

mimicking configurations of interacting amino-acid side chains as observed in a set of 47 high-resolution protein crystal structures. The energies correspond to the DW-CCSD(T<sup>\*\*</sup>)-F12 level of theory.<sup>41</sup>

#### D. Water cluster

We also include a dataset which we calculated for 4000 structures containing 40 water molecules, drawn from an NVE-molecular dynamics (MD) trajectory of a water droplet, simulated at 300 K, treated with the TIP3P potential<sup>42</sup> as implemented in CHARMM C41a1.<sup>43</sup> For each structure, a single-point energy was calculated at the DFT level using the PBEh-3c method.<sup>44</sup> Additional details, as well as the full dataset, can be found in the [supplementary material](#).

#### E. Solids: OQMD

We have used the Inorganic Crystal Structure Database<sup>45,46</sup> subset corresponding to the OQMD by Wolverton and co-workers.<sup>47,48</sup> This dataset has already been used to develop and benchmark random forest based QML model (Voronoi).<sup>49</sup> The dataset consists of ~30k crystal structures and formation energies, calculated using high-throughput DFT with generalized gradient approximation (GGA + U). We have used a random subset consisting of 3k structures with less than 40 atoms in the unit cell and formation energies lower than 5 eV/atom for training and testing. 1k crystals were used for testing and up to 2k crystals were used for training.

#### F. Solids: Elpasolites

We have also tested our representation for the elpasolite crystal structure dataset.<sup>19</sup> This dataset consists of ~10k elpasolite structures and DFT (PBE<sup>50</sup>) formation energies. The crystals correspond to quaternary main group elemental composition with all elements up to bismuth (39 in total). We have used a random subset of 7k structures, with up to 6k and 1k structures for training and testing, respectively.

#### G. Main group diatomics

To test the predictive power for alchemical interpolation, we have also included a set of previously published DFT (PBE<sup>50</sup>) results for single, double, and triple bonds among main-group diatomics saturated with hydrogens.<sup>17</sup> The training/test splits are generated differently, as explained in Sec. IV B.

### IV. RESULTS AND DISCUSSION

Using learning curves (generally resulting in straight lines when recorded on log-log plots due to their inverse power law relationship<sup>51</sup>), we first present numerical results which indicate the predictive power of our QML model for atomization and formation energies in various datasets. When available for the same dataset, we also compare to other QML models in the literature. Thereafter, the alchemical extrapolation capacity is demonstrated for predicting covalent bonds in molecules with elements that were not part of training. Finally,

log-log plots of learning curves for nine electronic ground-state properties of organic molecules (QM9) are reported and discussed.

#### A. Energies of molecules, clusters, and solids

Figure 6 displays the performance overview for energy predictions on six different datasets (QM9, QM7b, SSI, water, elpasolites, OQMD). Mean absolute out-of-sample energy prediction errors are shown as a function of training set size. The results indicate remarkable performance for all datasets, indicating a well-working QML model yielding systematic improvement with increasing training set size. The learning curves also indicate out-of-sample MAEs which are consistently lower, or similar, than previously published models in the literature. For QM9, the MAE reaches the highly coveted chemical accuracy threshold (1 kcal/mol or ~0.043 eV for enthalpy of formation) with only 2k training points on the QM9 dataset. Previously published QML models had to include an order of magnitude more training molecules to reach such accuracy. This is similar to the amount of training molecules necessary when using the Coulomb matrix representation in conjunction with semi-empirical or DFT based baselines in order to estimate electron correlated energies, as demonstrated in 2015 with the  $\Delta$ -ML model.<sup>52</sup>

For QM9, atomic Spectral London Axilrod-Teller-Muto (aSLATM)<sup>14</sup> and SOAP multi-kernel model<sup>20,53</sup> reach a performance nearly as good as our QML model. aSLATM, however, performs worse for the SSI and the water cluster. The SOAP multi-kernel QML model, however, performs an expansion in kernel function space acting on the distance for which all degrees of freedom have already been integrated out. As such it is, strictly speaking, not the same as an improved representation, but rather an improved regressor. Note that single kernel based SOAP QML models perform significantly worse. The reader should take notice however that in the SOAP learning curve results presented in Fig. 6, the ~3k structures which had failed the SMILES consistency test, were included. As such, these QML models are not exactly comparable, and the SOAP results are still likely to slightly improve if these faulty structures were to be removed. One should also note that the SOAP results shown for QM7b correspond to the multi-kernel SOAP kernel.<sup>20,54</sup>

Other models presented correspond to the Coulomb matrix (CM),<sup>7</sup> bags of bonds (BOB),<sup>8</sup> Bonds and Angles based Machine Learning (BAML),<sup>9</sup> Histogram of Distances, Angles, and Dihedrals (HDAD),<sup>3</sup> Spectral London Axilrod-Teller-Muto (SLATM), aSLATM,<sup>14</sup> crystal representation by Faber, Lindmaa, Lilienfeld, Armiento (FLLA),<sup>19</sup> sine matrix,<sup>55</sup> and many-body tensor representation (MBTR).<sup>15</sup> We also compared to QML models which are not based on KRR, such as the message passing neural network model (enns2s),<sup>5</sup> and a Voronoi-tessellation based random forest model (Voronoi).<sup>49</sup>

The MAE of our new QML model is consistently the lowest for all datasets and large training sets. For the set of 4000 non-equilibrium water clusters, there is a noticeable difference between the global (CM, BOB, and SLATM) and the atomic representations (i.e., aSLATM and the new model we introduce in this work): The global models exhibit very little



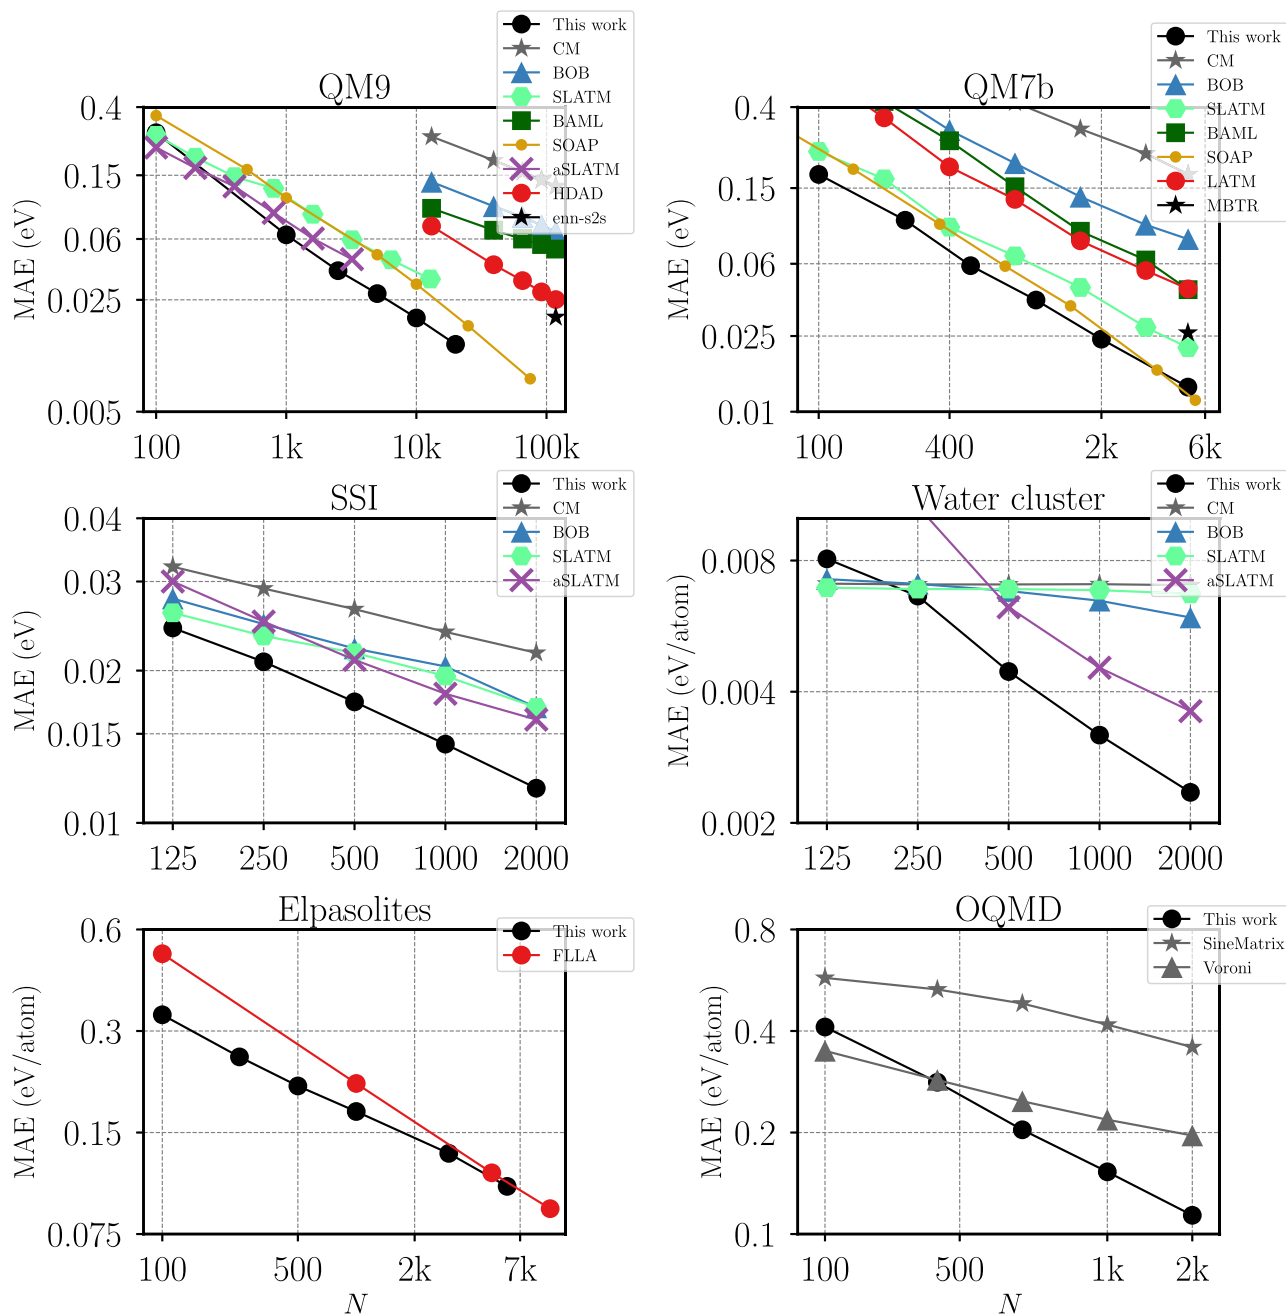


FIG. 6. Learning curves for atomization/formation energy predictions corresponding to various QML models. Out-of-sample MAE is shown as a function of training set size for molecules (QM9 and QM7b), protein side-chain dimers (SSI), liquid water ( $\text{H}_2\text{O}$ )<sub>40</sub> snapshots (Water cluster), and crystalline (OQMD and elpasolites) datasets.

learning at first, only for larger  $N$  the learning curves begin to turn downward. The atomic models, however, our new representation based QML model as well as aSLATM, improve rapidly with increasing training dataset size. We believe that sorting and crowding in the global representations make it difficult to accurately account for the purely geometrical changes in structures that contribute to total energy variations.

Impressive predictive power is also observed for the OQMD dataset, a structurally and compositionally very diverse set of solids. Our new model has a lower out-of-sample MAE for all  $N$  when compared to the sine matrix representation on the OQMD dataset. The offset of the learning curve of our new model is larger than that of the Voronoi-based

random-forest model.<sup>49</sup> However, the learning rate of our QML model is significantly steeper, surpassing the Voronoi model already at just  $\sim 250$  training samples. Results for a solid state variant of the CM, designed for use in periodic systems, have also been included (SineMatrix).<sup>55</sup> It has a similar slope as the Voronoi model, but a substantially larger offset.

For the elpasolite dataset,<sup>19</sup> with large composition diversity but identical crystal structures, the learning-curve of the FLLA representation has a slightly higher offset than our new QML model, yet exhibits a steeper learning curve, though our model converges toward the same slope for larger training set sizes. We can only speculate on the reasons for such behavior.

The FLLA representation differs qualitatively from the other representations in this study: It does not include any explicit information about coordinates and only encodes periodic row and column of the elements which populate each crystal structure site. The QML model then learns to infer ground state energies without knowing the exact configuration. This leads to a very low dimensional model that is still unique for the system, which might be the cause of the steep slope. This however needs to be investigated more carefully before any conclusions can be drawn.

## B. Alchemical predictions

Our new scaled many-body expansion explicitly accounts not only for distributions of interatomic distances and angles but also for elemental distributions in the periodic table. We have studied its capability to predict covalent binding of molecules containing chemical elements which were not present in the molecules used for training. More specifically, we have investigated single, double, and triple bonds with one bonding atom coming from group (IV), i.e., C, Si, or Ge. In order to increase covalent bond order, we have varied the valency of their bonding partner as follows: For single bonds, group IV atoms are bound to halogens (group VII). For double bonds, group IV atoms are bound to chalcogen atoms (group VI), and for triple bonds, group IV atoms are bound to group V atoms. Dangling valencies of group IV atoms have been saturated with hydrogen. Similar covalent bonding potentials have also recently been used in order to assess the predictive power of first- and second-order perturbation theory based alchemical predictions.<sup>17</sup>

In order to test the alchemical “extrapolation,” we trained on the covalent bonds of all other compounds (16 curves) which did contain neither the group IV atom nor the corresponding bonding partner in question. The predictive power for the out-of-sample molecule, as displayed in Fig. 7, is impressive. Albeit not quantitative (chemical accuracy is

not reached), the results are semi-quantitative and certainly provide a physically very adequate picture of the covalent bonding in single, double, and triple bonds for main-group atoms in periods 2–4. The fact that predictions for the central elements H<sub>2</sub>SiS are more accurate (easier to interpolate) than others is consistent with this interpretation. We also note that the deviation is the worst for 2nd-row elements (due to lack of *d*-orbitals they differ substantially more from 3rd and 4th rows than 3rd and 4th rows differ from each other). Because of poor performance, we do not compare to other representations in this test.

We have also investigated thousands of organic molecules in order to obtain improved quantitative statistics on the question of how out-of-sample prediction errors of different representations are affected when elements in the test set are excluded from training. We tested our new model, with and without elemental smearing ( $\sigma_P, \sigma_G = 1.6$  and  $\sigma_P, \sigma_G \rightarrow 0$ , respectively), and we also included BOB and CM representations for comparison. Details about how the training and test sets were selected can be found in Sec. III A.

Corresponding prediction errors [MAE, root mean squared error (RMSE), and maximal error] for test (training) sets with (without) nitrogen or with (without) oxygen are shown in Fig. 8. Obviously, the models perform best when both, training *and* test molecules, contain the same elements. However, even for models trained on sets without nitrogen/oxygen and without elemental smearing, considerable predictive power and, maybe more importantly, systematic improvement with training set size is observed for our new model. Use of elemental smearing results in a slight improvement.

The loss in accuracy due to the absence of elements in training is substantially worse for BOB and CM. Notably, BOB, in general considered to be more accurate than CM,<sup>3,8</sup> experiences a more dramatic loss than CM resulting in BOB being worse than CM. This can be understood if one considers the fact that BOB bags nuclear Coulomb repulsion

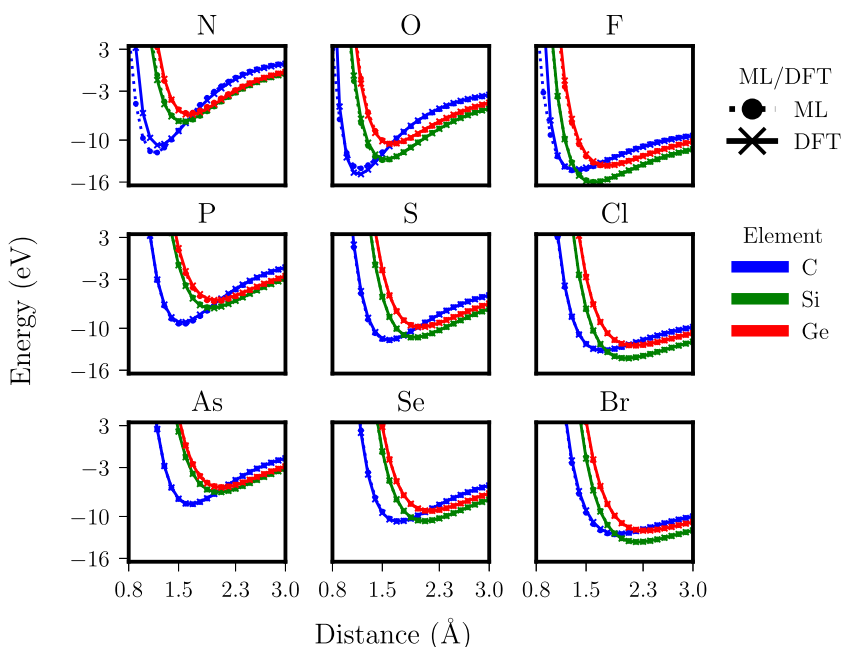


FIG. 7. Covalent bond potentials calculated by DFT (star) and estimated by QML (circle) for 27 main-group diatomic molecules. Bonding occurs between a group IV element (C: blue, Si: green, or Ge: red) and halogens (single bond), chalcogens (double bond), or a group V element (triple bond). Columns correspond to triple (left), double (middle), and single bonds (right). Rows correspond to the period of the group IV atom’s binding partner: 2nd period (top), 3rd period (middle), 4th period (bottom).

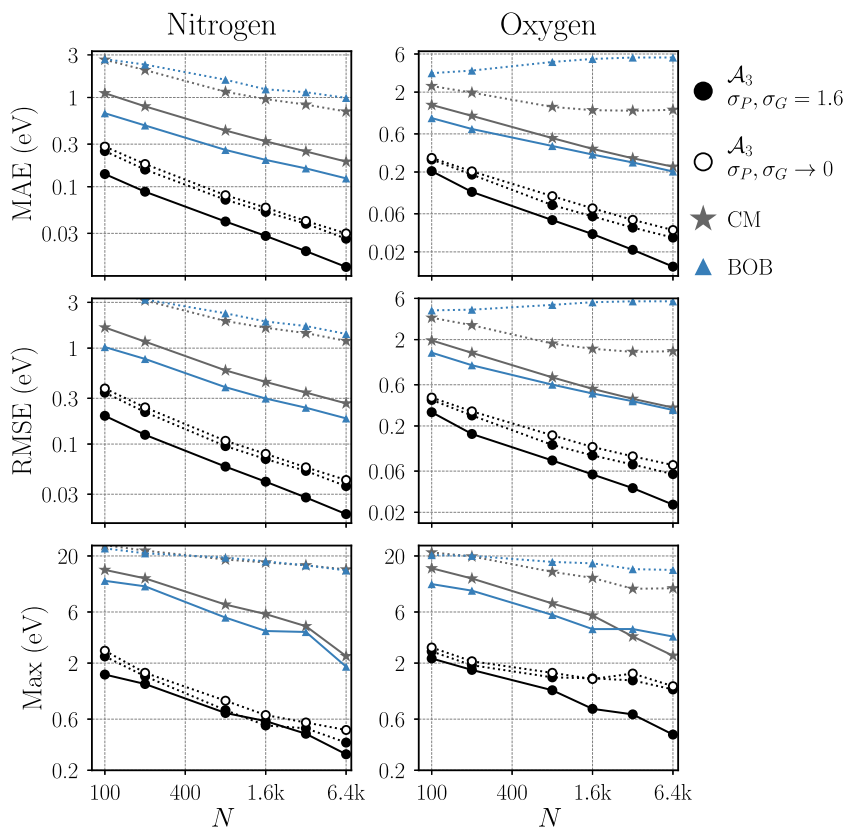


FIG. 8. Learning curves for atomization energy predictions using training sets with (full lines) and without (dotted lines) elements nitrogen (left) or oxygen (right). All molecules were extracted from QM9 (see Sec. III A). QML predictions have been made using our new model, with (filled circle) and without (empty circle) elemental smearing, as well as KRR with CM and BOB representation (following Refs. 8 and 56).

terms by element pairs, whereas the CM matrix directly compares the Coulomb interactions without (explicit) regard for which elements-pairs are being compared, effectively already performing an “alchemical” comparison. This allows the CM-based models to meaningfully interpolate even toward elements, not part of training. Our model, however, clearly outperforms CM and BOB: For example, its MAE reaches chemical accuracy ( $\sim 0.03$  eV) for the nitrogen lacking training set at  $\sim 6.4k$  molecules (rather than at  $\sim 1.6k$  training molecules containing N). Note that CM and BOB based KRR models are far from reaching such accuracy even after being trained on molecules *containing* the element in question!

The relatively high accuracy of our model achieved without any of the elemental smearing, however, is surprising. We would have expected that the lack of the appropriate elements in training introduces prediction errors which can no longer be decreased through the addition of more molecules. However, the learning curves in Fig. 8 do not indicate any worsening of the learning rate. While possible that the expected deterioration could still be observed for larger training sets, we do find it surprising that such high accuracy can be reached without any elemental smearing at all.

These results clearly demonstrate that alchemical extrapolation is possible when interpolating elemental groups and periods in the periodic table through an appropriate representation. Since the representation is continuous in the corresponding compositional space, we also believe that indication is given that the calculation of alchemical derivatives is meaningful, similar in spirit to Ref. 57.

### C. Other ground state properties of molecules

Finally, we also investigated how well QML models, based on our new representation and optimized for energies, perform when predicting other ground-state quantum properties, part of the QM9 dataset. More specifically, we have included atomization energies, highest occupied molecular orbital (HOMO) and lowest unoccupied molecular orbital (LUMO) eigenvalues as well as the HOMO-LUMO gap, dipole moment, polarizability, zero point vibrational energy (ZPVE), heat capacity, and the vibrational frequency of the highest lying fundamental ( $\omega$ ). Results are shown in Fig. 9 and provide overwhelming evidence that resulting models enable predictions systematically improving with training set size, no matter what property. For comparison, we have also included results for the aSLATM model. aSLATM results are typically worse when dealing with extensive properties, such as energies, polarizability, or heat-capacity. When dealing with intensive properties, such as eigenvalues or dipole moment, aSLATM is on par or even slightly better than our model, with the exception of  $\omega$ .  $\omega$  corresponds to the vibrational stretch frequency of CH, NH, or OH bonds, a property with hardly any variance at all. Previously we have seen that this property is best predicted by a random forest model which has relatively poor performance for all other properties.<sup>3</sup> The interpretation is that predicting this property is much more a classification problem than a supervised learning task.

Figure 9 also includes learning curves for the RMSE, indicating a slightly higher offset (to be expected) and systematic improvement with training set size with similar slopes as the

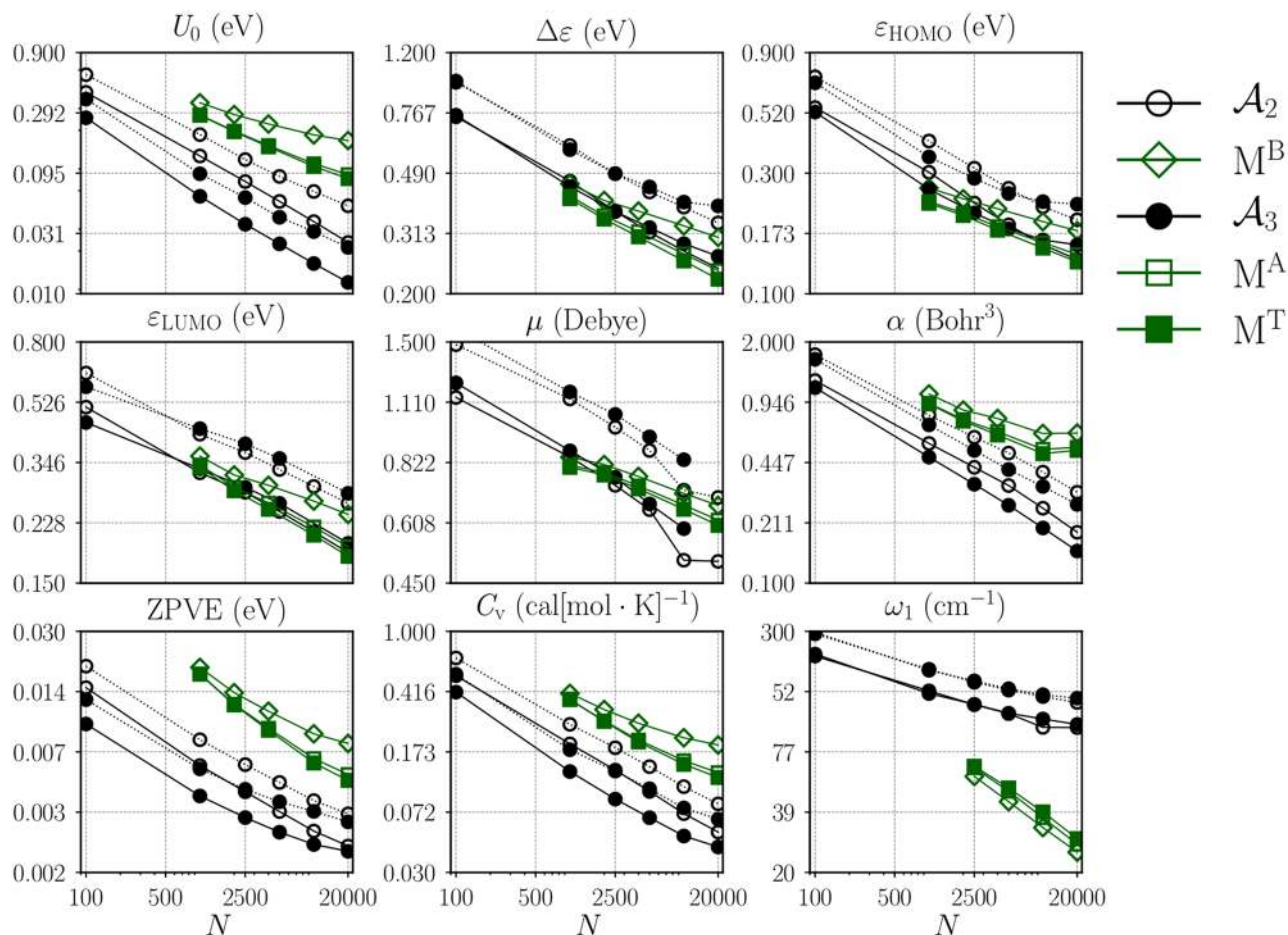


FIG. 9. Learning curves for out-of-sample MAE (filled lines) and RMSE (dashed lines) as a function of training set size  $N$  for nine electronic ground state properties in the QM9 dataset. QML predictions have been made using either a molecular kernel and BAML as representation or atomic kernels with our new representation. The BAML representation includes bonds ( $M^B$ ); bonds and angles ( $M^A$ ); and bonds, angles, and torsional angles ( $M^T$ ). Predicted properties include atomization energy, at 0 K ( $U_0$ ); HOMO-LUMO gap ( $\Delta\varepsilon$ ); HOMO eigenvalue ( $\varepsilon_{\text{HOMO}}$ ); LUMO eigenvalue ( $\varepsilon_{\text{LUMO}}$ ); norm of dipole moment ( $\mu$ ); static isotropic polarizability ( $\alpha$ ); zero point vibrational energy (ZPVE); heat capacity at room temperature ( $C_v$ ); and the highest fundamental vibrational frequency ( $\omega_1$ ).

mean absolute error. This is an assuring result, indicating that also predictions for outliers improve as training set size is increased, as already discussed in Refs. 9 and 52.

Furthermore, for Fig. 9 we have also distinguished between two- and three-body contributions (as well as four-body for BAML). For all properties but for  $\omega$ , the trend meets the expectation, as also already confirmed previously for BAML.<sup>9</sup> Addition of the higher order term systematically lowers the learning curves by a significant amount.

## V. CONCLUSION

We have introduced a universal representation of an atom in a chemical compound for use in QML models. An atom is represented by a sum of multidimensional Gaussians, each term corresponding to elemental, atom-pairwise, and angular distributions and scaled by respective power laws. For the compounds and properties studied, we have found four-body contributions to be insignificant. Most system-independent hyperparameters, such as exponents in scaling functions and Gaussian widths, were found to be not critical to the preference of resulting ML-models, as long as “reasonable” heuristics<sup>28</sup>

was used. This could, however, be explored more systematically within future work. Analytical expressions have been derived for corresponding distances between arbitrary chemical compounds. These distances can directly be used within KRR based QML models of electronic ground state properties. For energies of organic molecules, water clusters, amino-acid side chains, and crystalline solids, the resulting QML models lead to learning curves with a very low offset and steep learning rate. For compositionally diverse systems, chemical accuracy ( $\sim 1$  kcal/mol) can now be reached using only thousands of training instances. We have also studied the effect of explicitly accounting for inter-elemental distances in the periodic table: Our new QML model can produce semi-qualitatively accurate covalent bonding potentials for single, double, and triple bonds which include chemical element-pairs which were not part of training. For thousands of organic molecules, we also demonstrated that our model, after being trained on molecules which do not contain nitrogen or oxygen, still outperforms by a margin CM or BOB based models trained on molecules which do contain nitrogen or oxygen. For various electronic ground state properties of organic molecules, numerical results indicate that remarkable predictive power can be reached.

While the reference data used in this study has mostly been obtained at the hybrid DFT level of theory, the steep learning curves of our QML models suggest that it has now become a realistic possibility to obtain a sufficiently large training set at the post-Hartree-Fock level of theory (or from experiment) and to use it for the training of QML models which enable subsequent high-throughput screening efforts with similar accuracy.

Combining our new representation with the recently proposed QML model trained on molecular quasi-particles representing atoms-in-molecules (also known as “am-on” approach) might provide the possibility to generate accurate models which scale with the size of a query system.<sup>14</sup> Subsequent work will deal with forces and other properties.

## SUPPLEMENTARY MATERIAL

See [supplementary material](#) for a definition of the four body distribution ( $A_4$ ); derivation of the overlap integrals where the angular parts are solved in Fourier space; and a full description of the water cluster calculations as well as the water cluster data.

## ACKNOWLEDGMENTS

F.A.F. would like to thank Kuang-Yu Samuel Chang and Michele Ceriotti for data provided. O.A.v.L. acknowledges funding from the Swiss National Science Foundation (Nos. PP00P2\_138932 and 407540\_167186 NFP 75 Big Data). This research was partly supported by the NCCR MARVEL, funded by the Swiss National Science Foundation. Calculations were performed at sciCORE (<http://scicore.unibas.ch/>) scientific computing core facility at the University of Basel.

<sup>1</sup>F. Jensen, *Introduction to Computational Chemistry* (John Wiley, West Sussex, England, 2007).  
<sup>2</sup>F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, *Nat. Commun.* **8**, 872 (2017).  
<sup>3</sup>F. A. Faber, L. Hutchison, B. Huang, J. Gilmer, S. S. Schoenholz, G. E. Dahl, O. Vinyals, S. Kearnes, P. F. Riley, and O. A. von Lilienfeld, *J. Chem. Theory Comput.* **13**, 5255 (2017).  
<sup>4</sup>K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, *Nat. Commun.* **8**, 13890 (2017).  
<sup>5</sup>J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, in *Proceedings of the 34th International Conference on Machine Learning, ICML, 2017*.  
<sup>6</sup>D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, in *Advances in Neural Information Processing Systems (NIPS, 2015)*, pp. 2215–2223.  
<sup>7</sup>M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, *Phys. Rev. Lett.* **108**, 058301 (2012).  
<sup>8</sup>K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, *J. Phys. Chem. Lett.* **6**, 2326 (2015).  
<sup>9</sup>B. Huang and O. A. von Lilienfeld, *J. Chem. Phys.* **145**, 161102 (2016).  
<sup>10</sup>J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).  
<sup>11</sup>D. Rogers and M. Hahn, *J. Chem. Inf. Model.* **50**, 742 (2010).  
<sup>12</sup>O. A. von Lilienfeld, *Int. J. Quantum Chem.* **113**, 1676 (2013).  
<sup>13</sup>K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K. R. Müller, and E. K. U. Gross, *Phys. Rev. B* **89**, 205118 (2014).  
<sup>14</sup>B. Huang and O. A. von Lilienfeld, preprint [arXiv:1707.04146](https://arxiv.org/abs/1707.04146) (2017).  
<sup>15</sup>H. Huo and M. Rupp, preprint [arXiv:1704.06439](https://arxiv.org/abs/1704.06439) (2017).  
<sup>16</sup>O. A. von Lilienfeld, *Many-Electron Approaches in Physics, Chemistry and Mathematics* (Springer, 2014), pp. 169–189.  
<sup>17</sup>K. S. Chang, S. Fias, R. Ramakrishnan, and O. A. von Lilienfeld, *J. Chem. Phys.* **144**, 174110 (2016).  
<sup>18</sup>O. A. von Lilienfeld, R. Ramakrishnan, M. Rupp, and A. Knoll, *Int. J. Quantum Chem.* **115**, 1084 (2015).

<sup>19</sup>F. A. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *Phys. Rev. Lett.* **117**, 135502 (2016).  
<sup>20</sup>S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **18**, 13754 (2016).  
<sup>21</sup>K.-R. Müller, S. Mika, G. Rättsch, K. Tsuda, and B. Schölkopf, *IEEE Trans. Neural Networks* **12**, 181 (2001).  
<sup>22</sup>B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (MIT Press, 2002).  
<sup>23</sup>V. Vovk, “Kernel ridge regression,” in *Empirical Inference: Festschrift in Honor of Vladimir N. Vapnik*, edited by B. Schölkopf, Z. Luo, and V. Vovk (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013), pp. 105–116.  
<sup>24</sup>T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. (Springer, New York, 2011).  
<sup>25</sup>S. Mathias, M.Sc. thesis, University of Bonn, 2015, [http://wissrech.ins.uni-bonn.de/teaching/master/masterthesis\\_mathias\\_revised.pdf](http://wissrech.ins.uni-bonn.de/teaching/master/masterthesis_mathias_revised.pdf).  
<sup>26</sup>J. Barker, J. Bulin, J. Hamaekers, and S. Mathias, “LC-GAP: Localized Coulomb descriptors for the Gaussian approximation potential,” in *Scientific Computing and Algorithms in Industrial Simulations* (Springer, 2017).  
<sup>27</sup>A. P. Bartók and G. Csányi, *Int. J. Quantum Chem.* **115**, 1051 (2015).  
<sup>28</sup>R. Ramakrishnan and O. A. von Lilienfeld, *Chimia Int. J. Chem.* **69**, 182 (2015).  
<sup>29</sup>B. M. Axilrod and E. Teller, *J. Chem. Phys.* **11**, 299 (1943).  
<sup>30</sup>Y. Muto, *Proc. Phys. Math. Soc. Japan* **17**, 629 (1943).  
<sup>31</sup>A. S. Christensen, M. Elstner, and Q. Cui, *J. Chem. Phys.* **143**, 084123 (2015).  
<sup>32</sup>G. Pilania, J. E. Gubernatis, and T. Lookman, *Comput. Mater. Sci.* **129**, 156 (2017).  
<sup>33</sup>R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, *Sci. Data* **1**, 140022 (2014).  
<sup>34</sup>P. J. Stevens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, *J. Phys. Chem.* **98**, 11623 (1993).  
<sup>35</sup>L. Ruddigkeit, R. van Deursen, L. C. Blum, and J.-L. Reymond, *J. Chem. Inf. Model.* **52**, 2864 (2012).  
<sup>36</sup>G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, *New J. Phys.* **15**, 095003 (2013).  
<sup>37</sup>L. C. Blum and J.-L. Reymond, *J. Am. Chem. Soc.* **131**, 8732 (2009).  
<sup>38</sup>M. Ernzerhof and G. E. Scuseria, *J. Chem. Phys.* **110**, 5029 (1999).  
<sup>39</sup>C. Adamo and V. Barone, *J. Chem. Phys.* **110**, 6158 (1999).  
<sup>40</sup>L. A. Burns, J. C. Faver, Z. Zheng, M. S. Marshall, D. G. A. Smith, K. Vanommeslaeghe, A. D. MacKerell, Jr., K. M. Merz, Jr., and C. D. Sherrill, *J. Chem. Phys.* **147**, 161727 (2017).  
<sup>41</sup>M. S. Marshall and C. D. Sherrill, *J. Chem. Theory Comput.* **7**, 3978 (2011).  
<sup>42</sup>W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein, *J. Chem. Phys.* **79**, 926 (1983).  
<sup>43</sup>B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, and M. Karplus, *J. Comput. Chem.* **4**, 187 (1983).  
<sup>44</sup>S. Grimme, J. G. Brandenburg, C. Bannwarth, and A. Hansen, *J. Chem. Phys.* **143**, 054107 (2015).  
<sup>45</sup>A. Belsky, M. Hellenbrandt, V. L. Karen, and P. Luksch, *Acta Crystallogr., Sect. B: Struct. Sci.* **58**, 364 (2002).  
<sup>46</sup>G. Bergerhoff, R. Hundt, R. Sievers, and I. D. Brown, *J. Chem. Inf. Comput. Sci.* **23**, 66 (1983).  
<sup>47</sup>S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, *npj Comput. Mater.* **1**, 15010 (2015).  
<sup>48</sup>J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, *JOM* **65**, 1501 (2013).  
<sup>49</sup>L. Ward, R. Liu, A. Krishna, V. I. Hegde, A. Agrawal, A. Choudhary, and C. Wolverton, *Phys. Rev. B* **96**, 024104 (2017).  
<sup>50</sup>J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).  
<sup>51</sup>K.-R. Müller, M. Finke, N. Murata, K. Schulten, and S. Amari, *Neural Comput.* **8**, 1085 (1996).  
<sup>52</sup>R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, *J. Chem. Theory Comput.* **11**, 2087 (2015).  
<sup>53</sup>A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. Kermode, G. Csányi, and M. Ceriotti, *Sci. Adv.* **3**(12), e1701816 (2017).  
<sup>54</sup>A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).  
<sup>55</sup>F. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *Int. J. Quantum Chem.* **115**, 1094 (2015).  
<sup>56</sup>K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **9**, 3404 (2013).  
<sup>57</sup>O. A. von Lilienfeld, *J. Chem. Phys.* **131**, 164102 (2009).