





RESEARCH ARTICLE | AUGUST 02 2023

q-pac: A Python package for machine learned charge equilibration models

Special Collection: [Software for Atomistic Machine Learning](#)

Martin Vondrák ; Karsten Reuter ; Johannes T. Margraf  

 Check for updates

J. Chem. Phys. 159, 054109 (2023)

<https://doi.org/10.1063/5.0156290>



View
Online



Export
Citation

CrossMark

500 kHz or 8.5 GHz?
And all the ranges in between.

Lock-in Amplifiers for your periodic signal measurements



Find out more

 Zurich
Instruments

q-pac: A Python package for machine learned charge equilibration models

Cite as: J. Chem. Phys. 159, 054109 (2023); doi: 10.1063/5.0156290

Submitted: 28 April 2023 • Accepted: 18 July 2023 •

Published Online: 2 August 2023



View Online



Export Citation



CrossMark

Martin Vondrák,  Karsten Reuter,  and Johannes T. Margraf^{a)} 

AFFILIATIONS

Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, D-14195 Berlin, Germany

Note: This paper is part of the JCP Special Topic on Software for Atomistic Machine Learning.

^{a)} Author to whom correspondence should be addressed: margraf@fhi.mpg.de

ABSTRACT

Many state-of-the-art machine learning (ML) interatomic potentials are based on a local or semi-local (message-passing) representation of chemical environments. They, therefore, lack a description of long-range electrostatic interactions and non-local charge transfer. In this context, there has been much interest in developing ML-based charge equilibration models, which allow the rigorous calculation of long-range electrostatic interactions and the energetic response of molecules and materials to external fields. The recently reported kQEq method achieves this by predicting local atomic electronegativities using Kernel ML. This paper describes the q-pac Python package, which implements several algorithmic and methodological advances to kQEq and provides an extendable framework for the development of ML charge equilibration models.

© 2023 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/10.1063/5.0156290>

I. INTRODUCTION

Atomistic machine learning (ML) methods and interatomic potentials, in particular, have had an enormous impact in the fields of molecular and materials simulation.¹⁻⁵ One of the key innovations that made this possible was the idea of decomposing the total energy of a system into atomic contributions, which could be learned as a function of each atom's chemical environment within a certain cutoff radius using Neural Networks (NNs)⁶ or Gaussian Process Regression (GPR).⁷ This locality assumption has enabled the construction of highly accurate, computationally efficient, and size-extensive potentials that approach first-principles accuracy at a fraction of the cost.⁸⁻¹¹

At the same time, locality ultimately limits the achievable accuracy of a potential, since the information beyond the cutoff radius is not taken into account in this case.¹²⁻¹⁴ Indeed, long-range interactions can be substantial in bulk systems, most prominently due to the Coulomb interaction, which decays slowly ($\sim \frac{1}{r}$) with interatomic distance. These electrostatic interactions are often screened in practice so that local potentials can still effectively describe polar solids and liquids with surprising accuracy.^{9,10,14} Unfortunately, this cannot always be relied upon, however. For example, non-local charge

transfer at heterogeneous interfaces or through molecular wires cannot be adequately captured in this manner.¹⁵ Similarly, the relative stability of molecular crystal polymorphs sensitively depends on a balance of long-ranged electrostatic and dispersion interactions, precluding a purely local description.^{16,17}

Due to these limitations, the inclusion of long-range interactions in ML potentials has been an active field of study, with several different approaches in use. These, for example, include the use of global or non-local descriptors.¹⁸⁻²⁰ In many cases, physical baseline models can also provide the correct long-range physics at affordable computational cost (Δ -ML).^{16,17,21-23} Finally, message-passing neural networks can extend the range of local interatomic potentials by a multiple of the employed cutoff, although without including the full long-range interactions present in a periodic system.²⁴

In this contribution, we focus on approaches that tackle the problem of long-range electrostatics by describing the charge distribution of molecules or materials within the ML model itself (e.g., via partial charges). This has the advantage that it allows the incorporation of different total charge states and the response to external fields rigorously. The simplest approach to this end is to directly learn suitable reference charges, e.g., from Hirshfeld decomposition.^{25,26}

While this, in principle, affords a reasonable description of long-range electrostatics, it does not resolve the issue of non-local charge transfer since the charges are, in this case, themselves functions of a local ML model.

To overcome this, Goedecker and co-workers proposed the Charge Equilibration via Neural Network Technique (CENT),^{27–29} where the charges are obtained by minimizing a charge dependent energy function. Specifically, CENT uses the classical Charge Equilibration (QEq) model³⁰ as a basis, replacing fixed elemental electronegativities with environment-dependent ones predicted by a NN. This approach was subsequently developed further by Behler, Goedecker, and co-workers into the Fourth Generation High Dimensional Neural Network Potentials (4GHDNNP).^{1,15,31} Here, CENT and local NN potentials are combined, and partial charges are fitted to reproduce those obtained from Hirshfeld partitioning.³² Similarly, Xie *et al.* reported a self-consistent NN potential, where charges are obtained through the gradient-based minimization of a coupled local and electrostatic energy function.³³ Here, charges from the Becke population analysis were used as a reference for the partial charges.

Our recently reported Kernel Charge Equilibration (kQEq) method is in the same spirit as these approaches but uses Kernel ML instead of NNs.³⁴ Kernel methods are frequently used for interatomic potentials, as they are highly data-efficient, depend on few hyperparameters, and can be trained through a closed-form linear algebra expression.⁴ Furthermore, kQEq avoids the ambiguity of charge partitioning schemes by training directly on electrostatic observables, such as the dipole moment.

In this paper, we introduce the q-pac Python package. q-pac provides a modular framework for implementing machine-learned charge equilibration methods, with a particular focus on kQEq. We review the kQEq methodology and describe several new algorithmic and methodological advances in q-pac. In particular, the Kernel Ridge Regression (KRR) approach of the original kQEq paper is replaced by a sparse GPR formulation, which provides better computational scaling of training and prediction as a function of the training set size. Furthermore, additional fitting targets and the possibility to fit multiple properties at the same time have been implemented. Notably, this includes energies, which allow the development of fully long-ranged ML interatomic potentials based on kQEq. Finally, some example applications of these new capabilities are showcased.

II. THEORY

To provide a consistent account of the methodology, the kQEq working equations are rederived in this section, starting from the classical QEq approach of Rappe and Goddard.^{30,35} Differences and new features relative to the original implementation presented in Ref. 34 are highlighted where appropriate. Atomic units are used in all equations.

A. Charge equilibration

The core idea of QEq^{30,35} and related methods is to define a simple energy expression that depends on the charge distribution within a system. The ground-state charge distribution for a given geometry




FIG. 1. Illustration of the approximate electron density decomposition into an isolated atom density $\rho_0(\mathbf{r})$ (left) and fluctuation density $\delta\rho(\mathbf{r})$ (right) for a schematic one-dimensional acetylene molecule. The solid line represents the target electron density $\rho(\mathbf{r})$, the dotted line represents the superposition of isolated atom densities $\rho_0(\mathbf{r})$, and the dashed line represents a combination of $\rho_0(\mathbf{r})$ with an approximate fluctuation density $\delta\rho(\mathbf{r})$. The latter describes charge transfer and polarization within the molecule.

is then obtained by minimizing this energy, under the constraint that the total charge is conserved.

The charge distribution in a molecule or solid is rigorously described by the electron density $\rho(\mathbf{r})$ and the location of the nuclei. Since the electron density is a complex three dimensional distribution, it is computationally convenient to work with a more simplified representation, such as atomic partial charges, however. To this end, we can split the total electron density into a reference density $\rho_0(\mathbf{r})$ and a fluctuation term $\delta\rho(\mathbf{r})$, where the former is typically the superposition of electron densities of the corresponding isolated spherical atoms (see Fig. 1). Together with the corresponding nuclei, these atomic reference densities are charge neutral and, therefore, do not contribute to the long-range electrostatic interactions, leaving the fluctuation density as the object of interest.

In the following, we will assume some atomic partitioning of the fluctuation density,

$$\delta\rho(\mathbf{r}) = \sum_i \delta\rho_i(\mathbf{r}), \quad (1)$$

where $\delta\rho_i(\mathbf{r})$ is the local fluctuation density around atom i . This allows us to define partial charges as

$$q_i = \int \delta\rho_i(\mathbf{r}) d\mathbf{r}. \quad (2)$$

Note that since there is no unique partitioning of $\delta\rho(\mathbf{r})$, the partial charges are also to some extent arbitrary, although canonical choices like Hirshfeld partitioning exist.

We can now approximate the total energy of a non-periodic system as

$$E_{\text{tot}} \approx E_0 + E_{\text{QEq}} = E_0 + \underbrace{\sum_{i=1}^N \left(\chi_i q_i + \frac{1}{2} J_i q_i^2 \right)}_{\text{Site-Energy}} + \underbrace{\frac{1}{2} \iint \frac{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}d\mathbf{r}'}_{\text{Coulomb-Integral}}. \quad (3)$$

Here, the first term E_0 is a charge-independent reference energy. This could, e.g., be defined as the sum of the energies of the isolated

neutral atoms, or this could stem from a local (charge-independent) interatomic potential. The second term E_{QEq} collects all charge-dependent terms and will be our primary focus in the following. E_{QEq} can itself further be divided into two terms. The first of these is a site-energy term that sums over all N atoms i and represents the second-order Taylor expansion of the atomic energy with respect to the partial charges. In this context, the expansion coefficient χ_i is usually termed the electronegativity, while J_i is the electronic hardness. In the original QEq scheme, both of these coefficients are element-dependent parameters. The second term in E_{QEq} is the classical Coulomb energy of the fluctuation density.

In order to evaluate E_{QEq} , we now need to define a mathematical expression of the fluctuation density and its partitioning. Specifically, we will assume that the fluctuation density $\delta\rho(\mathbf{r})$ can approximately be expressed as a superposition of spherically symmetric atom-centered Gaussians. Each of these Gaussians is normalized to the corresponding atomic partial charge q_i according to Eq. (2) and has an inverse distribution width $\phi_i = 1/(\sqrt{2}\alpha_i)$, where α_i can be interpreted as an atomic radius. This leads to the expression

$$\delta\rho(\mathbf{r}) \approx \sum_{i=1}^N -q_i \left(\frac{\phi_i}{\sqrt{\pi}} \right)^3 \exp(-\phi_i^2 |\mathbf{r} - \mathbf{r}_i|^2). \quad (4)$$

Here, the Gaussians are centered at the atomic positions \mathbf{r}_i . Using this definition, the Coulomb integral in Eq. (3) can be evaluated analytically as

$$\iint \frac{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}d\mathbf{r}' = \sum_{i=1}^N \left(q_i^2 \frac{1}{2\alpha_i\sqrt{\pi}} + \sum_{j=1}^N q_i q_j \frac{\text{erf}\left(\frac{r_{ij}}{\sqrt{2}\gamma_{ij}}\right)}{r_{ij}} \right), \quad (5)$$

with $\gamma_{ij} = \sqrt{\alpha_i^2 + \alpha_j^2}$ and r_{ij} being the distance between atoms i and j . With this, E_{QEq} from Eq. (3) can be rewritten as

$$E_{\text{QEq}} = \sum_{i=1}^N \left[\chi_i q_i + \frac{1}{2} \left(J_i + \frac{1}{\alpha_i\sqrt{\pi}} \right) q_i^2 \right] + \frac{1}{2} \sum_{i,j}^N q_i q_j \frac{\text{erf}\left(\frac{r_{ij}}{\sqrt{2}\gamma_{ij}}\right)}{r_{ij}}. \quad (6)$$

Note that here the on-site contribution to the Coulomb integral has been pulled into the electronic hardness term. We can thus interpret the parameter J_i as a non-classical contribution to the hardness, while the classical contribution is given by the self-energy of the Gaussians.

In order to obtain the equilibrium partial charges q_i , E_{QEq} must now be minimized. Due to the chosen functional form of the site energy, this expression is quadratic in q_i , so that the optimal charges can be computed in closed form. To this end, we take the derivative of Eq. (6) with respect to each partial charge and set it to zero. This leads to the linear system of equations,

$$\frac{\partial E_{\text{QEq}}}{\partial q_i} = \sum_{j=1}^N A_{ij} q_j + \chi_i = 0, \quad (7)$$

with A_{ij} being the elements of the hardness matrix \mathbf{A} defined as

$$A_{ij} = \begin{cases} J_i + \frac{1}{\alpha_i\sqrt{\pi}} & \text{for } i = j, \\ \frac{\text{erf}\left(\frac{r_{ij}}{\sqrt{2}\gamma_{ij}}\right)}{r_{ij}} & \text{otherwise.} \end{cases} \quad (8)$$

Using a Lagrange multiplier λ to conserve the total charge q_{tot} , we obtain a linear system of equations that can be expressed in matrix notation as

$$\underbrace{\begin{pmatrix} A_{1,1} & A_{1,2} & \cdots & A_{1,N} & 1 \\ A_{2,1} & A_{2,2} & \cdots & A_{2,N} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ A_{N,1} & A_{N,2} & \cdots & A_{N,N} & 1 \\ 1 & 1 & \cdots & 1 & 0 \end{pmatrix}}_{\bar{\mathbf{A}}} \cdot \underbrace{\begin{pmatrix} q_1 \\ q_2 \\ \vdots \\ q_N \\ \lambda \end{pmatrix}}_{\bar{\mathbf{q}}} = - \underbrace{\begin{pmatrix} \chi_1 \\ \chi_2 \\ \vdots \\ \chi_N \\ -q_{\text{tot}} \end{pmatrix}}_{\bar{\chi}}. \quad (9)$$

Here, the bars denote that these arrays are expanded by one dimension due to the Lagrange multiplier. Without the bars, these symbols represent the corresponding N -dimensional arrays.

The charge vector $\bar{\mathbf{q}}$ (including the Lagrange multiplier λ) can now easily be computed as

$$\bar{\mathbf{q}} = -\bar{\mathbf{A}}^{-1} \bar{\chi}, \quad (10)$$

meaning that the charges are a linear function of the electronegativities. Using this matrix notation, E_{QEq} can be expressed as

$$E_{\text{QEq}} = \frac{1}{2} \bar{\mathbf{q}}^T \bar{\mathbf{A}} \bar{\mathbf{q}} + \bar{\mathbf{q}}^T \bar{\chi}. \quad (11)$$

B. Periodic boundary conditions

Up to this point (and in Ref. 34), we have only considered systems in open boundary conditions (i.e., isolated molecules in the gas phase). For periodic systems, Eqs. (10) and (11) can also be used. However, this requires using Ewald summation in the construction of the hardness matrix \mathbf{A} , in order to take the full long-range interactions in an infinite crystal into account.³⁶ In particular, the Coulomb integral must be modified. The same implementation as in 31 was adopted here.

Ewald summation allows the efficient computation of the electrostatic energy of N point charges in periodic boundary conditions by separating it into a real-space and a reciprocal-space contribution. To this end, each charge is embedded into an auxiliary Gaussian charge distribution of the opposite sign and width η , defined as³⁷

$$\eta = \frac{1}{\sqrt{2\pi}} V^{\frac{1}{3}}, \quad (12)$$

where V is the volume of the unit cell. Note that these auxiliary Gaussians are not to be confused with the ones defined in Eq. (4).

In the long-range, the electrostatic interactions between the point charges and the auxiliary charge distributions cancel out, so that the short-range part of the electrostatic energy E_{real} can be evaluated in real space as

$$E_{\text{real}} = \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^{N_{\text{neig}}} q_i q_j \frac{\text{erfc}\left(\frac{r_{ij}}{\sqrt{2}\eta}\right)}{r_{ij}}. \quad (13)$$

Here, the first sum goes over the N atoms i in the unit cell, and the second sum goes over all N_{neig} atoms j (including periodic replicas) within the cutoff distance r_{real} of atom i . The cutoff is derived from the width η of the auxiliary Gaussians and depends on the desired accuracy ϵ , which is a small positive number determined by the user (as in Ref. 31, we use a default value of $\epsilon = 10^{-8}$). This yields the following expression for r_{real} :^{37,38}

$$r_{\text{real}} = \sqrt{2}\eta\sqrt{-\log \epsilon}. \quad (14)$$

As a second step, the long-range interactions of the auxiliary charge distributions are computed. This can be evaluated efficiently in reciprocal space using the Fourier transform of the auxiliary charge density,

$$E_{\text{recip}} = \frac{2\pi}{V} \sum_{\mathbf{k} \neq 0} \frac{\exp\left(\frac{-\eta^2|\mathbf{k}|^2}{2}\right)}{|\mathbf{k}|^2} \left(\sum_{i=1}^N q_i \exp(i\mathbf{k} \cdot \mathbf{r}_i) \right)^2. \quad (15)$$

Here, the first sum goes over all reciprocal lattice points \mathbf{k} within the cutoff r_{recip} , which is computed as

$$r_{\text{recip}} = \frac{\sqrt{2}}{\eta} \sqrt{-\log \epsilon}. \quad (16)$$

The cutoff distance r_{recip} depends again on the user defined accuracy parameter ϵ .

Finally, the self-interaction of the auxiliary Gaussian charges is accounted for via

$$E_{\text{self}} = -\sum_{i=1}^N \frac{q_i^2}{\sqrt{2\pi}\eta}. \quad (17)$$

Summation of all previous terms is equal to the electrostatic energy of N point charges in periodic boundary conditions,

$$E_{\text{Ewald}} = E_{\text{real}} + E_{\text{recip}} + E_{\text{self}}. \quad (18)$$

Because we use Gaussian charge distributions of width α_i instead of point charges, an additional correction term is required³⁶

$$E_{\text{Gauss}} = -\frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^{N_{\text{neig}}} q_i q_j \frac{\text{erfc}\left(\frac{r_{ij}}{\sqrt{2}\eta}\right)}{r_{ij}} + \sum_{i=1}^N \frac{q_i^2}{2\sqrt{\pi}\alpha_i}. \quad (19)$$

Here, the first term is again applied for all interactions within the cutoff r_{real} , while the second term corresponds to the on-site

contribution of the Coulomb integral, which is also present in the non-periodic case.

The periodic Coulomb integral can now be written as

$$\iint \frac{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}d\mathbf{r}' = E_{\text{Ewald}} - \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i}^{N_{\text{neig}}} q_i q_j \frac{\text{erfc}\left(\frac{r_{ij}}{\sqrt{2}\eta}\right)}{r_{ij}} + \sum_{i=1}^N \frac{q_i^2}{2\sqrt{\pi}\alpha_i}. \quad (20)$$

From this, the periodic hardness matrix elements can be derived analogously to the non-periodic case.

C. Kernel charge equilibration

As described so far, conventional charge equilibration schemes like QEq require the definition of three parameters per element. These are the electronegativity (χ_i), the non-classical contribution to the hardness (J_i), and the atomic radius (α_i). In practice, this limits the achievable accuracy of QEq, since the same electronic properties are assumed for all atoms of the same element, independent of their chemical environment and oxidation state. ML based charge equilibration methods can overcome this limitation by allowing χ_i (and, in principle, also the other parameters) to adapt to the environment of each atom, e.g., via a NN.^{27,28}

To implement this environment dependence in a Kernel ML framework, kQEq expresses the electronegativities in terms of atomic environment representation vectors \mathbf{p}_i , a kernel function k , and regression weights w_m as

$$\chi_i(\mathbf{p}_i) = \sum_{m=1}^M k(\mathbf{p}_i, \mathbf{p}_m) w_m, \quad (21)$$

where the sum goes over all atoms m in a representative set of chemical environments. Simply put, this equation thus assigns the electronegativity of atom i based on the similarity between i and each atom m in the representative set, as quantified by the kernel function k .

In the original kQEq implementation reported in Ref. 34, the representative set simply consisted of all chemical environments in the training set. In this case, the cost of predicting the electronegativities scales linearly with the number of training samples. Even worse, the training cost of such a Kernel Ridge Regression (KRR) model scales cubically with the number of training samples. The new implementation in q-pac, therefore, uses a different regression framework, namely sparse GPR. This is directly analogous to the approach used in Gaussian Approximation Potentials (GAP).^{4,7} Specifically, the representative set now consists of M environments (also called sparse points), which form a representative subset of the training set. As in GAP, these are selected through a CUR decomposition of the matrix of representation vectors.^{4,39}

To represent the chemical environments of atoms, Smooth Overlap of Atomic Positions (SOAP)⁴⁰ vectors are used, as implemented in the Describe package. As in GAP, polynomial kernels are used to quantify similarities between SOAP vectors,

$$k(\mathbf{p}_i, \mathbf{p}_j) = (\mathbf{p}_i \cdot \mathbf{p}_j)^\zeta. \quad (22)$$

Throughout this manuscript, $\zeta = 2$ is used as a default. SOAP vectors are normalized so that $k(\mathbf{p}_i, \mathbf{p}_i) = 1$.

To predict the electronegativities of N atoms, Eq. (21) can be rewritten as a matrix–vector multiplication,

$$\boldsymbol{\chi} = \mathbf{K}_{NM} \mathbf{w}. \quad (23)$$

Here and in the following, we use the notation of Csányi and co-workers for Kernel matrices,⁴ where the subscripts indicate their dimensions. \mathbf{K}_{NM} is thus a matrix containing the evaluations of the kernel function between all M sparse points and all N environments to be predicted. Thus, the obtained electronegativities can then be used to predict charges via Eq. (10).

D. Training on electrostatic properties

In principle, Eqs. (23) and (10) fully specify the kQEq method. However, this leaves the key question of what the regression weights w_m should be. In Ref. 34, we showed that these can be computed in closed-form by solving a regularized least-squares problem. This was demonstrated by fitting kQEq models on molecular dipole moments. Importantly, the fact that training can be performed as a single closed-form linear algebra operation, in this case, hinges on the fact that dipole moments are linear functions of atomic partial charges. Additionally, the charges in QEq are linear functions of the electronegativities [see Eq. (10)], and the electronegativities are linear functions of the regression weights [see Eq. (23)]. Indeed, Kernel methods like KRR and GPR allow arbitrary linear transformations of the regression output in the loss function.

Taking advantage of this, q-pac provides a generalized loss function for fitting to any electrostatic property that is a linear function of atomic partial charges. The corresponding regularized least-squares loss reads

$$\begin{aligned} \mathcal{L}_t &= \|\mathbf{T}_t \bar{\mathbf{q}} - \mathbf{t}_{\text{ref}}\|_{\Sigma_t^{-1}}^2 + \|\mathbf{w}\|_{\mathbf{K}_{MM}}^2 \\ &= (\mathbf{T}_t \bar{\mathbf{q}} - \mathbf{t}_{\text{ref}})^T \Sigma_t^{-1} (\mathbf{T}_t \bar{\mathbf{q}} - \mathbf{t}_{\text{ref}}) + \mathbf{w}^T \mathbf{K}_{MM} \mathbf{w}, \end{aligned} \quad (24)$$

where \mathbf{t}_{ref} is a general target property (e.g., atomic charges or dipole vector elements), and \mathbf{T}_t is a transformation matrix that converts a Lagrange multiplier expanded vector [see Eq. (9)] of charges $\bar{\mathbf{q}}$ to the target property. Σ_t is an N -dimensional diagonal regularization matrix that contains noise parameters σ_i^2 , which are proportional to the regularization strength. Unlike the unit-less regularization parameter in the previous KRR implementation, σ_i has the unit of the predicted property and can be interpreted as the expected accuracy of the fit. Furthermore, the GPR framework allows assigning individual values of σ_i in order to weight training samples differently. For simplicity, a single value of σ_i is used for each of the examples below.

We now aim to find the vector of weights \mathbf{w} , which minimizes this loss function. To this end, Eq. (24) must be rewritten so that it only depends on \mathbf{w} . Here, a technical difficulty arises, in that Eq. (23) yields a vector of electronegativities $\boldsymbol{\chi}$, while Eq. (10) requires the extended vector $\tilde{\boldsymbol{\chi}}$, which includes the total charge of the system q_{tot} . This transformation is achieved via an auxiliary matrix \mathbf{X} ,

$$\begin{aligned} \underbrace{\begin{pmatrix} \chi_1 \\ \chi_2 \\ \vdots \\ \chi_N \\ -q_{\text{tot}} \end{pmatrix}}_{\tilde{\boldsymbol{\chi}}} &= \underbrace{\begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \\ 0 & 0 & \cdots & 0 \end{pmatrix}}_{\mathbf{X}} \cdot \underbrace{\begin{pmatrix} K_{1,1} & K_{1,2} & \cdots & K_{1,M} \\ K_{2,1} & K_{2,2} & \cdots & K_{2,M} \\ \vdots & \vdots & \ddots & \vdots \\ K_{N,1} & K_{N,2} & \cdots & K_{N,M} \end{pmatrix}}_{\mathbf{K}_{NM}} \\ &\cdot \underbrace{\begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_M \end{pmatrix}}_{\mathbf{w}} - \underbrace{\begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ q_{\text{tot}} \end{pmatrix}}_{\mathbf{q}_{\text{tot}}}. \end{aligned} \quad (25)$$

Plugging this into Eq. (10) yields

$$\bar{\mathbf{q}} = -\bar{\mathbf{A}}^{-1} \tilde{\boldsymbol{\chi}} = -\bar{\mathbf{A}}^{-1} \underbrace{(\mathbf{X} \mathbf{K}_{NM} \mathbf{w} - \mathbf{q}_{\text{tot}})}_{\mathbf{x}}. \quad (26)$$

Finally, the transformation matrix \mathbf{T}_t determines the targeted property \mathbf{t} ,

$$\mathbf{t} = \mathbf{T}_t \bar{\mathbf{q}} = -\mathbf{T}_t \bar{\mathbf{A}}^{-1} (\mathbf{X} \mathbf{K}_{NM} \mathbf{w} - \mathbf{q}_{\text{tot}}). \quad (27)$$

In the current implementation, transformation matrices for charges and dipoles are provided,

$$\mathbf{T}_q = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}, \quad (28)$$

$$\mathbf{T}_\mu = \begin{pmatrix} x_1 & x_2 & \cdots & x_N & 0 \\ y_1 & y_2 & \cdots & y_N & 0 \\ z_1 & z_2 & \cdots & z_N & 0 \end{pmatrix}. \quad (29)$$

Note that these transformation matrices can, in principle, easily be modified to accommodate for higher multipole moments or other charge-derived electrostatic properties, such as electrostatic potentials at given grid points.

The final form of the loss function is obtained by plugging Eq. (27) into Eq. (24), yielding

$$\begin{aligned} \mathcal{L}_t &= \|\mathbf{t} - \mathbf{t}_{\text{ref}}\|_{\Sigma_t^{-1}}^2 + \|\mathbf{w}\|_{\mathbf{K}_{MM}}^2 \\ &= \|\mathbf{T}_t \bar{\mathbf{A}}^{-1} (\mathbf{X} \mathbf{K}_{NM} \mathbf{w} - \mathbf{q}_{\text{tot}}) - \mathbf{t}_{\text{ref}}\|_{\Sigma_t^{-1}}^2 + \|\mathbf{w}\|_{\mathbf{K}_{MM}}^2. \end{aligned} \quad (30)$$

In this loss function, both the least-squares and regularization terms are quadratic in the regression weights \mathbf{w} . The optimal weights are

obtained by taking the derivative $\nabla_{\mathbf{w}}\mathcal{L}$ (which is linear in \mathbf{w}), setting it to zero and solving for \mathbf{w} .

While Eq. (30) is a general loss function for single-property prediction, it may also be of interest to train models on multiple properties simultaneously. To this end, q-pac also allows combined loss functions. For instance, one may want to fit a model that reproduces molecular dipoles with partial charges that are close to some population analysis scheme,

$$\mathcal{L}_{q/\mu} = \|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{ref}}\|_{\Sigma_{\mu}^{-1}}^2 + \|\mathbf{q} - \mathbf{q}_{\text{ref}}\|_{\Sigma_q^{-1}}^2 + \|\mathbf{w}\|_{\mathbf{K}_{MM}}^2. \quad (31)$$

Here, separate regularization parameters can be used for the different properties. This allows weighting of the properties relative to each other.

Note that for simplicity, all expressions provided herein assume a single kQEq problem with N atoms (i.e., one simulation cell or molecule). In practice, models are trained on multiple systems simultaneously using blocked matrices and concatenated vectors, which then naturally allow the use of systems with varying numbers of atoms N in the training set.

E. Training on energies

As discussed in the introduction, one of the main motivations for developing ML-based charge equilibration models is the development of interatomic potentials with full long-range electrostatics. To this end, training on reference charges (e.g., from Hirshfeld partitioning) can yield a reasonable description of long-range interactions. However, population analysis schemes are, in general, not optimal for this purpose, as the charges usually yield quantitatively incorrect electrostatic properties. More critically, the energy E_{QEq} is in our experience rather unphysical when only training on charges or dipole moments. This is due to the fact that the energy expression is only a latent quantity in this case (yielding the appropriate charges upon minimization), which bears no relation to the real potential energy surface. As a consequence, the on-site energy contributions can be large and overly sensitive to small geometric changes, making E_{QEq} a poor basis for an interatomic potential.

One way to overcome this issue is to simply ignore the site-energy term in the interatomic potential. This is the approach taken in the 4GHDNNPs mentioned above.¹⁵ However, this has the downside that the corresponding potentials are not self-consistent, in the sense that their charges do not minimize the energy. The other alternative is to explicitly include energies in the loss function so that E_{QEq} takes physical information about the potential energy surface into account.

Unfortunately, fitting energies is not entirely straightforward within the kQEq framework. This is because E_{QEq} is not linear in the charges, which means that a closed-form solution for the optimal regression weights does not exist. However, for a given set of charges, the energy is linear in the electronegativities. q-pac, therefore, includes a form of fixed-point iteration to obtain accurate self-consistent energy models.

Specifically, we use an arbitrary set of initial charges \mathbf{q}^0 (e.g., from Hirshfeld partitioning) and train electronegativities χ^1 that yield optimal charges for these charges. Subsequently, we predict the self-consistent charges \mathbf{q}^1 corresponding to χ^1 . In general, there is a large difference between \mathbf{q}^0 and \mathbf{q}^1 , so the self-consistent energies of this kQEq model will be rather inaccurate. However, iteratively

restarting this process usually yields significant improvements, in that the self-consistent charges \mathbf{q}^t corresponding to χ^t quickly converge toward the ones used to fit the energies (\mathbf{q}^{t-1}). In pathological cases, the loss function can be expanded to include a bias toward the charges from the previous iteration, further aiding convergence. Here, a practical approach is to begin training on energies alone until the energy fitting error starts to increase. The charge bias can then be added with an initially large regularization parameter $\sigma_q^2 = 0.01 e$ (corresponding to a small weight of charges in the loss function), which is subsequently decreased by a factor of 0.5 at each iteration.

By monitoring the energy fitting error and the charge differences between two iterations, optimal regression weights for a kQEq interatomic potential can be selected. The typical convergence of charges and energies in this process is illustrated for a set of ZnO nanoparticles in Fig. 2 (see below for details on the dataset). This shows that even without the charge bias, energies and charges converge well, with the differences between self-consistent and fitting charges being below 10^{-3} electron charges.

F. Atomic forces

To apply interatomic potentials in molecular dynamics (MD) simulations and geometry relaxations, it is essential to efficiently obtain energy derivatives with respect to atomic positions \mathbf{r}_j . While the derivative of Eq. (6) with respect to \mathbf{r}_j is straightforward to compute, a complication arises because both the charges \mathbf{q} and the electronegativities χ also depend on \mathbf{r}_j . Here, the use of self-consistent charges is beneficial because by definition,

$$\frac{\partial E_{\text{QEq}}}{\partial q_i} = 0. \quad (32)$$

Consequently, the force on atom j can be expressed as

$$\mathbf{F}_j = -\sum_{i=1}^N \left(q_i \frac{\partial \chi_i}{\partial \mathbf{r}_j} \right) + \sum_{i>j}^N q_i q_j \frac{\partial V_{ij}}{\partial \mathbf{r}_j}, \quad (33)$$

with

$$V_{ij} = \frac{\text{erf}\left(\frac{r_{ij}}{\sqrt{2}\gamma_{ij}}\right)}{r_{ij}}. \quad (34)$$

Here, the first term describes the force caused by the response of the atomic electronegativities to changes in \mathbf{r}_j . According to Eq. (21), this term only requires taking the derivatives of the kernel function and the SOAP vectors, the latter of which are obtained through Describe.

Meanwhile, the second term is simply the derivative of the shielded Coulomb energy, which reads

$$\frac{\partial V_{ij}}{\partial \mathbf{r}_j} = \frac{\sqrt{2}\gamma_{ij} \exp\left(-\frac{r_{ij}^2}{2\gamma_{ij}}\right) - \sqrt{\pi} \text{erf}\left(\frac{r_{ij}}{\sqrt{2}\gamma_{ij}}\right)}{\sqrt{\pi} r_{ij}^3}. \quad (35)$$

G. Technical aspects

q-pac is implemented as an object-oriented library using Python 3.9. It heavily relies on numpy⁴¹ for array operations and




FIG. 2. RMSE of the training set (blue) and validation set (red) energies as a function of fixed point iterations for a set of ZnO nanoparticles (left). RMSE between self-consistent charges and charges from the previous fixed point iteration (right). See Sec. III for a description of the ZnO nanoparticles and the employed training and validation sets.

linear algebra, `ase`^{42,43} for representing structural data and running atomistic simulations, and `Describe`³⁷ for calculating SOAP vectors and their derivatives. The Ewald summation portion of the code is written in C++. C++ types are exposed to Python via `pybind11`.⁴⁴

III. RESULTS

A. Effect of sparsification

The main algorithmic advance in `q-pac` relative to the original `kQEq` implementation is the use of sparse GPR instead of KRR for predicting electronegativities. In order to demonstrate the benefit of this change, a series of dipole moment prediction models were trained, similar to Ref. 34. Specifically, 35 000 randomly selected molecules from the QM9 database were used, spanning a variety of small organic molecules containing C, H, O, N, and F.⁴⁵ The corresponding reference dipole moments were computed at the PBE0/def2-TZVP level using ORCA.^{46–48} From this, 1000 molecules each were randomly drawn as validation and test sets, while differently sized training sets were randomly drawn from the remaining 33 000 molecules.

As described in Sec. III, multiple hyperparameters need to be defined for any `kQEq` model. Following Ref. 34, the non-classical atomic hardness J_i was set to 0 for all elements. Atomic radii α_i for all elements are tabulated in the original `QEq` paper.³⁰ In our previous work, we found it beneficial to scale these since they are not necessarily ideal for the Gaussian charge distributions used in `q-pac`. In this paper, all radii are globally scaled by $\frac{1}{\sqrt{2}}$, which yielded robust models in all cases we considered.

The regularization strength σ_μ was optimized for each training set using grid search on the validation set error. The main SOAP hyperparameter to be chosen is the cut-off radius r_{cut} , which was set to 4.4 Å. The remaining SOAP hyperparameters are discussed in the Supplementary Information, together with results for smaller values of r_{cut} .

To establish the accuracy of the sparse GPR approach, Fig. 3 shows the mean absolute error (MAE) in predicted dipole moments for different training sets as a function of the number of sparse points M (per element). For comparison, the dashed lines in these figures show results for the corresponding full GPR models, where M includes all chemical environments in the training set. This is

equivalent to the KRR models used in the original `kQEq` paper.³⁴ For training sets of 3000 and 10 000 molecules, the accuracy of the full GPR is reached with $M = 3000$ (0.12 and 0.1 D, respectively). Note that for the training set of 20 000 molecules, the full model was not computed due to prohibitively high memory requirements. With sparse GPR, this is not an issue, however, allowing even lower MAEs.

In Fig. 4, the evolution of test set errors and computational costs for training and prediction with the number of training molecules is shown. Here, it should be emphasized that the number of sparse points is given in terms of chemical environments, whereas the training set size is given in terms of the number of molecules. Depending on the training set size, sparse models with 1000 and 3000 sparse points are equally accurate as the full models, with the $M = 1000$ models deviating from the full results for training sets larger than 1000 molecules. As mentioned above, the full models become computationally prohibitive for the largest training set of 20 000 molecules.

This can also be seen from the timings in Fig. 4, which clearly show that the complete training process (including evaluation of the validation set for hyperparameter tuning) scales much more steeply

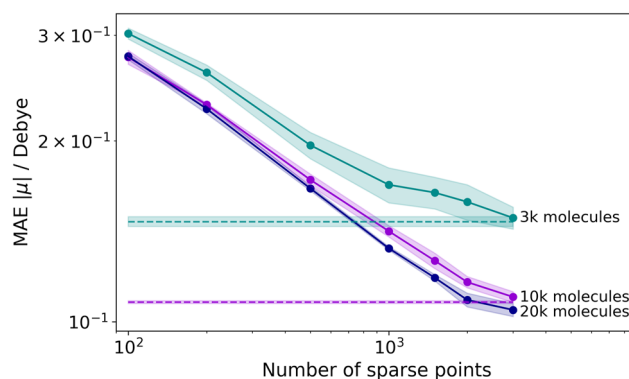


FIG. 3. Comparison of sparse (full line) and full (dashed line) `kQEq` models for training sets of 3000 (cyan), 10 000 (purple), and 20 000 (blue) molecules. Three different randomized training sets are used and averaged for each point.




FIG. 4. Top: Learning curves for dipole moment prediction using sparse ($M = 1000$ and $M = 3000$) and full ($M = N_{\text{train}}$) GPR. Bottom: Time needed for a full training cycle including CUR decomposition (where applicable), model training and validation (solid lines, left y-axis), and timings for 1000 dipole moment predictions (dashed lines, right y-axis). These calculations were performed on nodes with 2 Intel(R) Xeon(R) Platinum 8360Y CPUs (2.4 GHz) and 2048 GB RAM.

with the training set size for the full model. This is in line with the expected asymptotic scalings of $\mathcal{O}(N^3)$ and $\mathcal{O}(N)$ for the full and sparse models, respectively, although the sparse models are not yet in the linear regime in this plot. In concrete terms, training and validation of the full models took an average of 2.1 h (see supplementary material for hardware details) for 10 000 training molecules. In comparison, the $M = 3000$ model was six times faster, reaching nearly identical accuracy in 0.35 h. With 1000 sparse points per element, the same results were produced in 0.24 h.

The difference between sparse and full GPR models is even more striking when looking at the prediction times. Here, the sparse models scale as $\mathcal{O}(1)$, while the full models scale as $\mathcal{O}(N)$. This clearly has significant implications for models with large training sets and/or a large number of required predictions, where sparse kQEq models can potentially be one to two orders of magnitude more efficient. Beyond this substantial acceleration, the memory requirements of training sparse models are also much smaller.

B. Interatomic potentials for isolated systems

To illustrate the capabilities of kQEq for fitting potential energy surfaces (PES) of ionic materials, we developed an ML potential for a set of ZnO nanoparticles taken from the global optimization study of Chen and Dixon.⁴⁹ Specifically, a set of 98 low-energy structures of sizes between 62 and 264 zinc and oxygen atoms was used. Reference energies and Hirshfeld charges were computed with FHI-Aims at the PBE level using *tight* basis sets and integration settings.^{50,51}

Note that the choice of a predominantly ionic material like ZnO allows fitting accurate interatomic potentials using kQEq alone. Although bonding in ZnO is not purely ionic but contains a partial covalent character (in addition to other interactions like dispersion and exchange repulsion), the use of environment-dependent electronegativities allows kQEq models to indirectly take such effects into account. This works well for systems where non-ionic interactions are not dominating and of low body-order. In the more general case, such interactions should be treated separately.

The original dataset of Chen and Dixon exclusively consists of locally relaxed structures. In order to also test kQEq potentials on non-equilibrium structures, active learning was used to augment the dataset.^{4,52} Specifically, the initial kQEq interatomic potential (trained on the relaxed structures) was used to generate new configurations via MD simulations. From these, a diverse subset was selected via CUR decomposition, evaluated by the reference DFT method, and added to the training set. This active learning loop was repeated until the energy RMSE on newly generated structures converged.

The corresponding MD simulations were run through the Atomic Simulation Environment (ASE)^{42,43} calculator implemented in q-pac. Langevin dynamics were performed at 300 K with a 0.5 fs time step and a friction coefficient of 0.01. Before each production run, structures were reoptimized with the BFGS algorithm, followed by a 1 ps equilibration run. To account for the increasing accuracy of the interatomic potentials as a function of the active learning iterations, the length of the production runs was incrementally increased. Specifically, the initial simulation time was set to 0.2 ps and increased by a factor of two in each following iteration. Accordingly, the number of configurations added to the training set was also increased in each iteration, starting with 50 configurations.

In terms of model hyperparameters, the regularization for the energy term was set to $\sigma_E^2 = 0.01$ eV throughout. As discussed in the methods section, a charge bias term was added in later fixed point iterations to ensure convergence, although this was usually not necessary (see Fig. 2). The hardness parameters J_i were set to 27.21 eV (1.0 Ha) for Zn and O, and the SOAP cutoff was set to 3.0 Å. A full table of hyperparameters is provided in the supplementary material. Note that for this proof-of-principle application, no hyperparameter optimization was performed. It is, therefore, likely that even better performance could be achieved in principle.

The results of the active learning iterations are shown in Fig. 5. Here, the energy RMSE for new MD configurations is shown for




FIG. 5. Active learning potential for ZnO nanoparticles. The RMSE of predicted energies for new MD configurations generated with the potential at each active learning iteration is shown. For each iteration, the number of (added) training configurations is shown. The final RMSE was computed for 200 unseen configurations.

each iteration, along with the number of configurations added to the training set. The RMSE quickly drops from 39.6 meV/atom for the initial model to 1 meV/atom in iteration 2 and 0.8 meV/atom in the final iteration.

Beyond this good energetic accuracy, it is also of interest to consider the charge distributions that are learned by this potential. In Fig. 6 (top), the correlation between Hirshfeld and kQEq charges is shown. This reveals that kQEq charges are somewhat larger than Hirshfeld ones, with average charges of ± 0.54 and ± 0.37 , respectively. This is consistent with the known tendency of

Hirshfeld charges to underestimate charge transfer in polar systems, which is related to the use of neutral isolated atom densities to define the partitioning.^{53,54}

Notably, the variation of the charges displays an inverse correlation between kQEq and Hirshfeld, in particular for the oxygen atoms. Here, the least negative atoms in kQEq are the most negative in Hirshfeld and vice versa. This can be understood by considering the individual cases displayed in the lower part of Fig. 6, which reveals that Hirshfeld charges tend to be more extreme for undercoordinated edge and corner atoms, while kQEq charges are

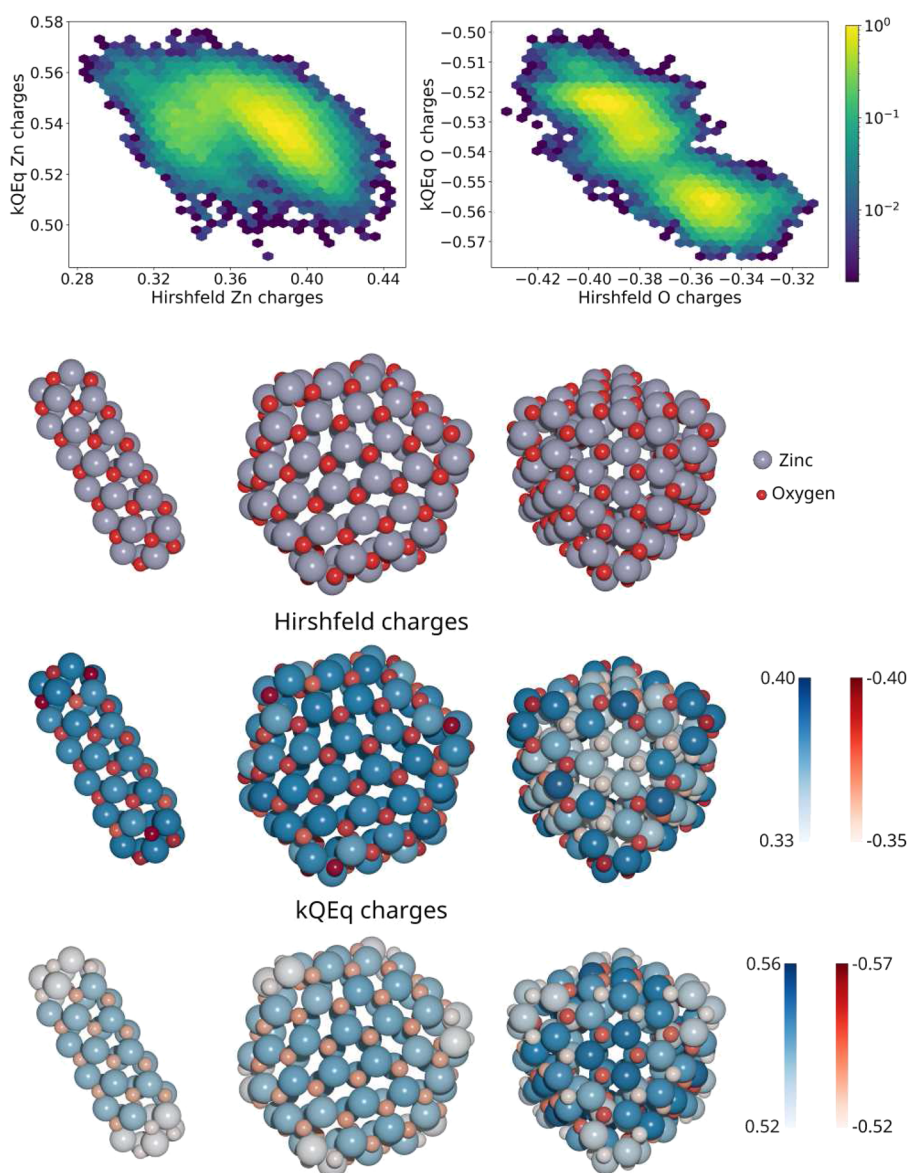


FIG. 6. Top: Correlation between Hirshfeld and kQEq charges for ZnO nanoparticles. Bottom: Comparison of charges for three representative structures, namely a nanotube, a hollow sphere, and a dense particle. Note that different color bars are used for Hirshfeld and kQEq charges since the latter are significantly larger in absolute terms.

less extreme for these atoms. This is again consistent with the way the respective charges are calculated. In Hirshfeld population analysis, the electron density is spatially partitioned according to a stockholder scheme where the electron densities of isolated neutral atoms are used to define the respective weights. Here, lower coordination environments mean that the density is partitioned between fewer partners. Meanwhile, kQEq charges minimize an electrostatic energy expression. Here, more extreme charges can be stabilized through electrostatic interactions in higher coordination environments.

Overall, partial charges are, in general, somewhat arbitrary, so no strong conclusions in favor of either model can be drawn from the charges alone. Here, we can consider dipole moments as an observable instead. While dipole interactions are, in principle, included in energies, they only dominate in the long-range interactions between clusters. Since we only train on isolated clusters, it should not be expected that these effects are learned by the current model. Indeed, while kQEq dipole elements correlate reasonably well with the DFT reference, they deviate rather strongly in absolute terms (see Fig. S6). Notably, the Hirshfeld charges obtained from DFT calculations actually reproduce the dipole of the particles surprisingly well in this case, as shown in Fig. S7. This indicates that combined loss functions with energies, charges, and/or dipole moments could be used to obtain improved electrostatic properties. Importantly, the kQEq charges are optimized for producing an interatomic potential for these clusters, whereas the Hirshfeld charges are not.

So far, we have only mentioned the non-classical atomic hardness J as a hyperparameter without discussing its effects in detail. Qualitatively speaking, larger values of J have the effect of counteracting charge transfer. To demonstrate how this affects the predicted charges and energies, 25 additional models were trained on the ZnO particle dataset, with different combinations of hardness parameters for Zn and O (see Fig. S8). This reveals that higher values of J lead to smaller absolute values of the charges and also decrease the fluctuations for the individual charge types. This is also reflected in the correlations between DFT and kQEq dipole moments (see Fig. S6), where models using smaller hardness parameters are better at reproducing the DFT dipole moments, likely because smaller hardness allows charges to be distributed more freely. Meanwhile, the energy RMSEs change less than 0.3 meV per atom between the best and worst models, underscoring that the energy is rather insensitive to this parameter.

C. Interatomic potentials for periodic systems

As described in the methods section, q-pac also allows the development of models for periodic systems. To demonstrate this, we fit an interatomic potential for a range of bulk structures with the stoichiometries ZnO and ZnO₂.

An initial set of crystals was obtained from Materials Project⁵⁵ (16 for ZnO and 9 for ZnO₂, see supplementary material). This set was augmented by creating random neutral vacancy defect pairs (i.e., removing one O and one Zn atom) in different supercells of each crystal, yielding 20 configurations for each supercell. Here, the supercells were used to sample different defect densities and were chosen in order to still allow reference DFT calculations for all cells (i.e., containing less than 400 atoms, see supplementary material

for details). Additionally, non-equilibrium structures were generated by randomly perturbing the atomic positions with Gaussian noise, yielding 20 additional structures per supercell (half of them perturbed with $\sigma = 0.05$ Å, the other half with $\sigma = 0.1$ Å). This led to a total of 1025 structures with supercells ranging from 52 to 384 atoms, for which reference calculations were performed at the PBE level using FHI-Aims with *light* basis set and integration settings.

In order to train the kQEq model, a training set was generated from this data by randomly drawing 150 configurations each for pristine and defective ZnO structures, respectively. Similarly, 80 ZnO₂ configurations each were drawn from the pristine and defective sets. The remaining 606 structures were used as an unseen test set. The same hyperparameters as for the non-periodic structures were used for the kQEq and SOAP settings.

The corresponding model displays an RMSE of 4.1 meV/atom on the test set. This is quite satisfactory, given that the test set covers an energy span of nearly 1 eV/atom (see Fig. 7). Importantly, the potential is able to fit ZnO and ZnO₂ on the same footing, since it is able to describe different oxidation states of Zn through environment dependent electronegativities.

This is also reflected in Fig. 8, which shows the correlation between Hirshfeld and kQEq charges for the test set. Both approaches yield distinct clusters for the ZnO and ZnO₂ charges, with the charges again being larger in magnitude for kQEq (e.g., ± 0.55 vs ± 0.35 for ZnO). Notably, the Hirshfeld oxygen charges display a strikingly large variation for ZnO₂ and are even positive in some cases, whereas the corresponding kQEq charges are consistently negative and display much smaller variance.

To illustrate the corresponding charge distributions, three simulation cells for a tetraauricupride-structured ZnO polymorph (MP-ID 13161) are shown in Fig. 8. As a reference, Hirshfeld and kQEq charge distributions for the relaxed supercell are shown in the left column. In the center, a rattled configuration is shown. kQEq and Hirshfeld charges display a qualitatively similar response to this perturbation. Finally, the right frame shows a structure with a vacancy pair. Here, an O atom is missing on the bottom right, and a Zn atom is missing on the bottom left. Interestingly, the response to the O vacancy is almost identical between the kQEq and

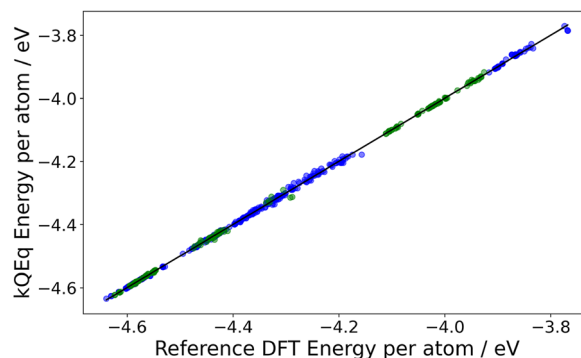


FIG. 7. Parity plot of kQEq predicted energies and reference DFT calculations for a test set of periodic ZnO (blue) and ZnO₂ (green) structures. The root mean squared error of the model is 4.1 meV/atom.

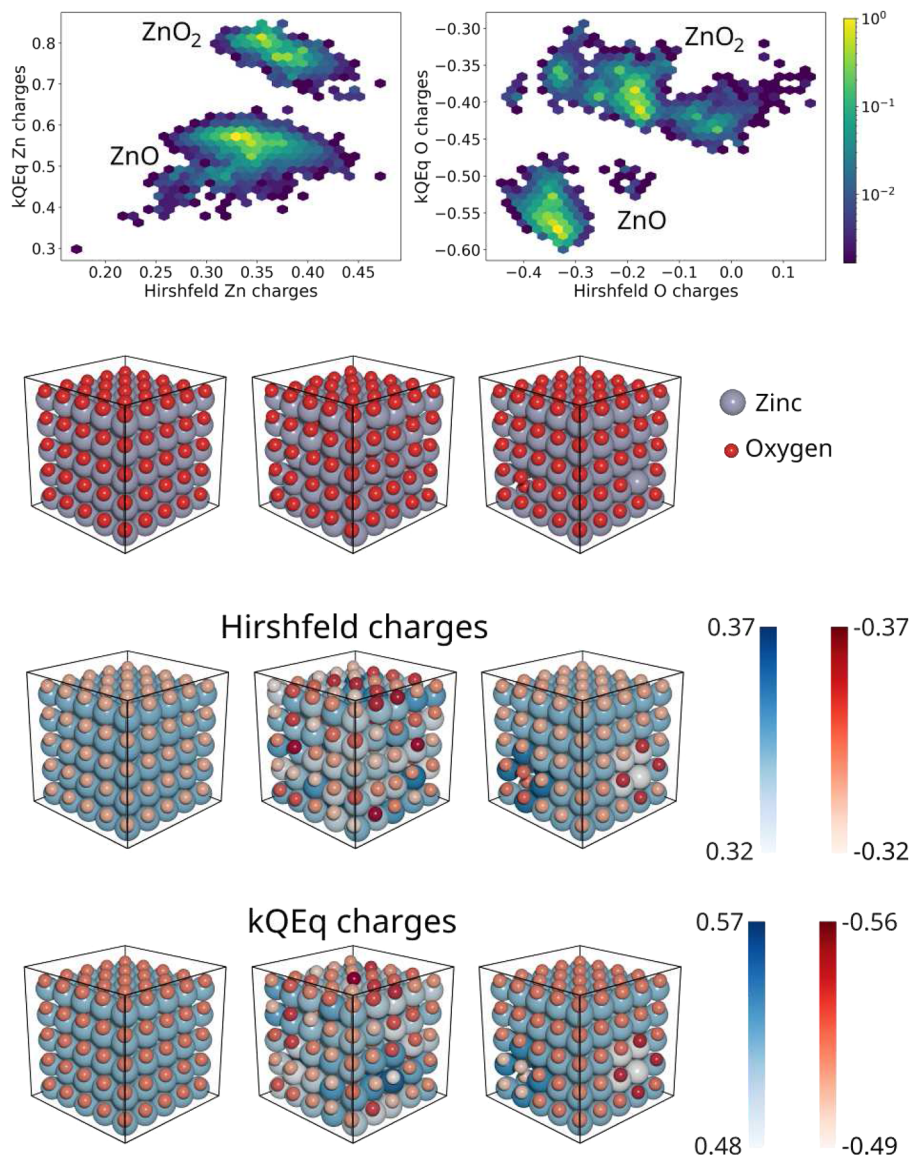


FIG. 8. Top: Correlation between Hirshfeld and kQE charges for ZnO and ZnO₂ bulk structures. Bottom: Comparison of charges for a relaxed ZnO structure, randomly rattled structure, and structure with a vacancy pair (MP ID 13161, $5 \times 5 \times 5$ supercells). Note that different color bars are used for Hirshfeld and kQE charges since the latter are significantly larger in absolute terms.

Hirshfeld charges, with the adjacent O atoms being more negative and the adjacent Zn atoms being less positive, consistent with an overall charge neutral defect. Meanwhile, the response for the Zn vacancy is similar for the adjacent Zn atoms (which become more positive) but different for the adjacent oxygen atoms. Here, the corresponding Hirshfeld charges are slightly more negative than average, while the kQE charges are slightly less negative. Nevertheless, the agreement between these methods in terms of the electronic localization of the defect is rather good, given that both methods are based on very different premises.

Since we present models for zinc oxides targeting either isolated clusters or periodic structures, it is natural to ask whether the models are transferable from one case to the other. This is not quite straightforward to test for the reported models, since the bulk structures feature both ZnO and ZnO₂ compositions, whereas the clusters all display the stoichiometry ZnO. Additionally, slightly different DFT settings were used to train both models.

To explore the transferability between bulk and cluster models, we, therefore, retrained the periodic model using only ZnO compositions and consistent DFT settings with the isolated case.

Additionally, a new model was trained on both bulk and cluster configurations. In Fig. S5, the transferability is shown for these three models. This reveals that the predictions of the bulk model are systematically too high in energy and display significant scatter for cluster configurations. In the reverse case, even stronger deviations are observed, with some configurations being predicted to be several eV too high in energy. Nonetheless, when training a model on both types of configurations, excellent predictive performance is achieved on the test set. This indicates that the lack of transferability is not due to the simultaneous inclusion of periodic and cluster configurations but rather due to inherent differences in the two sets. For example, the bulk set contains defective and randomly perturbed structures, while the clusters feature hollow nanospheres and tubes with atomic environments that strongly differ from the bulk case.

D. Forces

As for other interatomic potentials, in principle, forces can also be used as fitting targets for kQEq. Currently, this feature is not implemented in q-pac, however. This is mainly due to the increased memory demand when fitting on forces as opposed to energies. We are actively working to overcome this limitation at the moment. Nevertheless, it is already of interest to consider how accurate forces are for kQEq models trained on energies.

As described above, the model for ZnO particles was trained on molecular dynamics configurations in an active learning setup. Due to these non-equilibrium structures, the model has sufficient information about the shape of the potential energy surface to predict forces with reasonable accuracy (RMSE = 0.17 eV/Å, see Fig. S3). In contrast, the potential for bulk ZnO and ZnO₂ structures is less suitable for use in dynamical calculations due to overall larger force errors (see Fig. S4). Here, the training configurations were created by randomly removing atoms from crystalline supercells (without further relaxation) or randomly perturbing atomic positions. Both of these procedures lead to rather unphysical structures and large forces. In practice, active learning would also allow better performance on force prediction for the bulk structures, however.

IV. CONCLUSIONS

In this paper, we have introduced the q-pac package, which provides an efficient and general framework for fitting kQEq models. This is achieved through a new sparse GPR formulation of the kQEq method and the implementation of additional and generalizable fitting targets. Importantly, this allows fitting kQEq interatomic potentials for the first time using a fixed point iteration algorithm.

While the showcased applications demonstrate the functionality and accuracy of this approach, pure kQEq potentials are (much like the CENT method)²⁷ limited to predominantly ionic materials like ZnO. In future work, we will explore hybrid potentials that combine short-ranged local interatomic potentials like GAP with kQEq, in order to obtain generally applicable potentials with full long-range interactions.

The modularity of q-pac will also allow for the development of ML charge equilibration models beyond the simple QEq energy

function. This may become necessary for properly modeling the response of polarizable systems to external fields, where conventional QEq is known to be inadequate.⁵⁶ However, there are some indications that the much larger flexibility of kQEq mitigates these problems to a large extent.³⁴ Either way, q-pac provides the ideal testing ground for addressing such pathologies both in terms of the physics of charge equilibration and in terms of more advanced ML approaches.

SUPPLEMENTARY MATERIAL

See the supplementary material for hyperparameters, a detailed description of ZnO data, and an additional analysis of predicted charges and forces.

ACKNOWLEDGMENTS

We thank Jonas Finkler for his help with implementing electrostatic interactions for periodic systems.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Martin Vondrák: Investigation (equal); Methodology (equal); Software (lead); Visualization (lead); Writing – original draft (equal). **Karsten Reuter:** Conceptualization (equal); Funding acquisition (lead); Writing – review & editing (equal). **Johannes T. Margraf:** Conceptualization (lead); Methodology (equal); Software (supporting); Supervision (lead); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The code and data for this paper are publicly available at <https://gitlab.com/jmargraf/qpac>.

REFERENCES

- ¹J. Behler, “Four generations of high-dimensional neural network potentials,” *Chem. Rev.* **121**, 10037–10072 (2021).
- ²J. T. Margraf, H. Jung, C. Scheurer, and K. Reuter, “Exploring catalytic reaction networks with machine learning,” *Nat. Catal.* **6**, 112–121 (2023).
- ³O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, and K.-R. Müller, “Machine learning force fields,” *Chem. Rev.* **121**, 10142–10186 (2021).
- ⁴V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, and G. Csányi, “Gaussian process regression for materials and molecules,” *Chem. Rev.* **121**, 10073–10141 (2021).
- ⁵J. Margraf, “Science-driven atomistic machine learning,” *Angew. Chem., Int. Ed.* **62**, e202219170 (2023).
- ⁶J. Behler and M. Parrinello, “Generalized neural-network representation of high-dimensional potential-energy surfaces,” *Phys. Rev. Lett.* **98**, 583 (2007).
- ⁷A. 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).

- ⁸S. Stocker, J. Gasteiger, F. Becker, S. Günemann, and J. T. Margraf, “How robust are modern graph neural network potentials in long and hot molecular dynamics simulations?,” *Mach. Learn.: Sci. Technol.* **3**, 045010 (2022).
- ⁹V. Kapil, C. Schran, A. Zen, J. Chen, C. J. Pickard, and A. Michaelides, “The first-principles phase diagram of monolayer nanoconfined water,” *Nature* **609**, 512–516 (2022).
- ¹⁰T. Morawietz, A. Singraber, C. Dellago, and J. Behler, “How van der waals interactions determine the unique properties of water,” *Proc. Natl. Acad. Sci. U. S. A.* **113**, 8368–8373 (2016).
- ¹¹V. L. Deringer, N. Bernstein, G. Csányi, C. Ben Mahmoud, M. Ceriotti, M. Wilson, D. A. Drabold, and S. R. Elliott, “Origins of structural and electronic transitions in disordered silicon,” *Nature* **589**, 59–64 (2021).
- ¹²M. Herbold and J. Behler, “A hessian-based assessment of atomic forces for training machine learning interatomic potentials,” *J. Chem. Phys.* **156**, 114106 (2022).
- ¹³V. L. Deringer and G. Csányi, “Machine learning based interatomic potential for amorphous carbon,” *Phys. Rev. B* **95**, 094203 (2017).
- ¹⁴C. G. Staacke, H. H. Heenen, C. Scheurer, G. Csányi, K. Reuter, and J. T. Margraf, “On the role of long-range electrostatics in machine-learned interatomic potentials for complex battery materials,” *ACS Appl. Energy Mater.* **4**, 12562–12569 (2021).
- ¹⁵T. W. Ko, J. A. Finkler, S. Goedecker, and J. Behler, “A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer,” *Nat. Commun.* **12**, 398 (2021).
- ¹⁶S. Wengert, G. Csányi, K. Reuter, and J. T. Margraf, “Data-efficient machine learning for molecular crystal structure prediction,” *Chem. Sci.* **12**, 4536–4546 (2021).
- ¹⁷S. Wengert, G. Csányi, K. Reuter, and J. T. Margraf, “A hybrid machine learning approach for structure stability prediction in molecular co-crystal screenings,” *J. Chem. Theory Comput.* **18**, 4586–4593 (2022).
- ¹⁸H. Jung, S. Stocker, C. Kunkel, H. Oberhofer, B. Han, K. Reuter, and J. T. Margraf, “Size-extensive molecular machine learning with global representations,” *ChemSystemsChem* **2**, 659 (2020).
- ¹⁹S. Chmiela, V. Vassilev-Galindo, O. T. Unke, A. Kabylda, H. E. Sauceda, A. Tkatchenko, and K.-R. Müller, “Accurate global machine learning force fields for molecules with hundreds of atoms,” *Sci. Adv.* **9**, 1875 (2023).
- ²⁰A. Grisafi and M. Ceriotti, “Incorporating long-range physics in atomic-scale machine learning,” *J. Chem. Phys.* **151**, 204105 (2019).
- 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**, 2087–2096 (2015).
- ²²G. Fan, A. McSloy, B. Aradi, C.-Y. Yam, and T. Frauenheim, “Obtaining electronic properties of molecules through combining density functional tight binding with machine learning,” *J. Phys. Chem. Lett.* **13**, 10132–10139 (2022).
- ²³J. Westermayr, S. Chaudhuri, A. Jeindl, O. T. Hofmann, and R. J. Maurer, “Long-range dispersion-inclusive machine learning potentials for structure search and optimization of hybrid organic–inorganic interfaces,” *Digital Discovery* **1**, 463–475 (2022).
- ²⁴K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, “Quantum-chemical insights from deep tensor neural networks,” *Nat. Commun.* **8**, 190 (2017).
- ²⁵N. Artrith, T. Morawietz, and J. Behler, “High-dimensional neural-network potentials for multicomponent systems: Applications to Zinc Oxide,” *Phys. Rev. B* **83**, 153101 (2011).
- ²⁶T. Morawietz, V. Sharma, and J. Behler, “A neural network potential-energy surface for the water dimer based on environment-dependent atomic energies and charges,” *J. Chem. Phys.* **136**, 064103 (2012).
- ²⁷S. A. Ghasemi, A. Hofstetter, S. Saha, and S. Goedecker, “Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network,” *Phys. Rev. B* **92**, 045131 (2015).
- ²⁸S. Faraji, S. A. Ghasemi, S. Rostami, R. Rasoulkhani, B. Schaefer, S. Goedecker, and M. Amsler, “High accuracy and transferability of a neural network potential through charge equilibration for calcium fluoride,” *Phys. Rev. B* **95**, 1041 (2017).
- ²⁹E. R. Khajepasha, J. A. Finkler, T. D. Kühne, and S. A. Ghasemi, “CENT2: Improved charge equilibration via neural network technique,” *Phys. Rev. B* **105**, 144106 (2022).
- ³⁰A. K. Rappe and W. A. I. Goddard, “Charge equilibration for molecular dynamics simulations,” *J. Phys. Chem.* **95**, 3358–3363 (1991).
- ³¹T. W. Ko, J. A. Finkler, S. Goedecker, and J. Behler, “General-purpose machine learning potentials capturing nonlocal charge transfer,” *Acc. Chem. Res.* **54**, 808–817 (2021).
- ³²F. L. Hirshfeld, “Bonded-atom fragments for describing molecular charge densities,” *Theor. Chim. Acta* **44**, 129–138 (1977).
- ³³X. Xie, K. A. Persson, and D. W. Small, “Incorporating electronic information into machine learning potential energy surfaces via approaching the ground-state electronic energy as a function of atom-based electronic populations,” *J. Chem. Theory Comput.* **16**, 4256–4270 (2020).
- ³⁴C. G. Staacke, S. Wengert, C. Kunkel, G. Csányi, K. Reuter, and J. T. Margraf, “Kernel charge equilibration: Efficient and accurate prediction of molecular dipole moments with a machine-learning enhanced electron density model,” *Mach. Learn.: Sci. Technol.* **3**, 015032 (2022).
- ³⁵S. Ramachandran, T. G. Lenz, W. M. Skiff, and A. K. Rappé, “Toward an understanding of Zeolite Y as a cracking catalyst with the use of periodic charge equilibration,” *J. Phys. Chem.* **100**, 5898–5907 (1996).
- ³⁶T. R. Gingrich and M. Wilson, “On the Ewald summation of Gaussian charges for the simulation of metallic surfaces,” *Chem. Phys. Lett.* **500**, 178–183 (2010).
- ³⁷L. Himanen, M. O. Jäger, E. V. Morooka, F. Federici Canova, Y. S. Ranawat, D. Z. Gao, P. Rinke, and A. S. Foster, “Dscribe: Library of descriptors for machine learning in materials science,” *Comput. Phys. Commun.* **247**, 106949 (2020).
- ³⁸R. A. Jackson and C. R. A. Catlow, “Computer simulation studies of zeolite structure,” *Mol. Simul.* **1**, 207–224 (1988).
- ³⁹M. W. Mahoney and P. Drineas, “CUR matrix decompositions for improved data analysis,” *Proc. Natl. Acad. Sci. U. S. A.* **106**, 697–702 (2008).
- ⁴⁰A. P. Bartók, R. Kondor, and G. Csányi, “On representing chemical environments,” *Phys. Rev. B* **87**, 184115 (2013).
- ⁴¹C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, “Array programming with NumPy,” *Nature* **585**, 357–362 (2020).
- ⁴²A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dulak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rossgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, “The atomic simulation environment—A python library for working with atoms,” *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- ⁴³S. R. Bahn and K. W. Jacobsen, “An object-oriented scripting interface to a legacy electronic structure code,” *Comput. Sci. Eng.* **4**, 56–66 (2002).
- ⁴⁴W. Jakob, J. Rhineland, and D. Moldovan, “pybind11 – seamless operability between C++11 and Python,” (2017), <https://github.com/pybind/pybind11>.
- ⁴⁵R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, “Quantum chemistry structures and properties of 134 kilo molecules,” *Sci. Data* **1**, 140022 (2014).
- ⁴⁶F. Neese, “The orca program system,” *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2**, 73–78 (2012).
- ⁴⁷F. Neese, “Software update: The orca program system – version 5.0,” *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **12**, e1606 (2022).
- ⁴⁸J. P. Perdew, M. Ernzerhof, and K. Burke, “Rationale for mixing exact exchange with density functional approximations,” *J. Chem. Phys.* **105**, 9982–9985 (1996).
- ⁴⁹M. Chen and D. A. Dixon, “Machine-learning approach for the development of structure–energy relationships of ZnO nanoparticles,” *J. Phys. Chem. C* **122**, 18621–18639 (2018).
- ⁵⁰V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, “Ab initio molecular simulations with numeric atom-centered orbitals,” *Comput. Phys. Commun.* **180**, 2175–2196 (2009).

⁵¹J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized gradient approximation made simple,” *Phys. Rev. Lett.* **77**, 3865–3868 (1996).

⁵²N. Artrith and J. Behler, “High-dimensional neural network potentials for metal surfaces: A prototype study for copper,” *Phys. Rev. B* **85**, 045439 (2012).

⁵³P. Bultinck, C. Van Alsenoy, P. W. Ayers, and R. Carbó-Dorca, “Critical analysis and extension of the Hirshfeld atoms in molecules,” *J. Chem. Phys.* **126**, 144111 (2007).

⁵⁴A. V. Marenich, S. V. Jerome, C. J. Cramer, and D. G. Truhlar, “Charge model 5: An extension of Hirshfeld population analysis for the accurate description

of molecular interactions in gaseous and condensed phases,” *J. Chem. Theory Comput.* **8**, 527–541 (2012).

⁵⁵A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, “Commentary: The materials project: A materials genome approach to accelerating materials innovation,” *APL Mater.* **1**, 011002 (2013).

⁵⁶J. P. Koski, S. G. Moore, R. C. Clay, K. A. O’Hearn, H. M. Aktulga, M. A. Wilson, J. A. Rackers, J. M. D. Lane, and N. A. Modine, “Water in an external electric field: Comparing charge distribution methods using ReaxFF simulations,” *J. Chem. Theory Comput.* **18**, 580–594 (2022).