



PAPER

OPEN ACCESS

RECEIVED
24 November 2021REVISED
14 February 2022ACCEPTED FOR PUBLICATION
18 February 2022PUBLISHED
11 March 2022

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Kernel charge equilibration: efficient and accurate prediction of molecular dipole moments with a machine-learning enhanced electron density model

Carsten G Staacke¹, Simon Wengert¹, Christian Kunkel¹, Gábor Csányi², Karsten Reuter¹
and Johannes T Margraf^{1,*}

¹ Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, D-14195 Berlin, Germany

² Engineering Laboratory, University of Cambridge, Cambridge, CB2 1PZ, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: margraf@fhi.mpg.de

Keywords: kernel, charge equilibration, partial charges, dipole moment

Supplementary material for this article is available [online](#)

Abstract

State-of-the-art machine learning (ML) interatomic potentials use local representations of atomic environments to ensure linear scaling and size-extensivity. This implies a neglect of long-range interactions, most prominently related to electrostatics. To overcome this limitation, we herein present a ML framework for predicting charge distributions and their interactions termed kernel charge equilibration (kQEq). This model is based on classical charge equilibration (QEq) models expanded with an environment-dependent electronegativity. In contrast to previously reported neural network models with a similar concept, kQEq takes advantage of the linearity of both QEq and Kernel Ridge Regression to obtain a closed-form linear algebra expression for training the models. Furthermore, we avoid the ambiguity of charge partitioning schemes by using dipole moments as reference data. As a first application, we show that kQEq can be used to generate accurate and highly data-efficient models for molecular dipole moments.

1. Introduction

Kernel and neural network (NN) based machine-learning (ML) methods have in recent years become established as an essential addition to the toolbox of computational chemistry [1–4]. In particular, ML-based interatomic potentials have had great success in providing energies and forces with quantum mechanical accuracy at a fraction of the cost of first-principles calculations [1, 5–10]. To achieve size-extensivity and a linear computational scaling with system size, these ML potentials typically rely on a local representation of atomic environments and consequently assume that the energy can be decomposed into local atomic contributions [9, 11]. This simple idea has led to a strong focus of chemical ML research on developing sophisticated representations of local atomic environments and, relatedly, NN architectures that directly embed atoms in their neighborhood [11–14, 14–18].

At the same time, it is clear that the assumption of locality does not hold for all systems to the same extent [19]. Indeed, strongly polar or ionic systems display very long-ranged Coulomb interactions. Even for a fairly unpolar (e.g. organic) system, the locality of the energy does not necessarily imply that other electronic properties are similarly local. In particular, electronic properties such as molecular orbital energies or dipole moments can break locality assumptions [20]. Consequently, such properties tend to be more challenging to predict with purely data-driven ML methods [16, 20]. Beyond this methodological challenge, dipole moments are an intrinsically interesting target as they govern the asymptotic decay of interactions of neutral molecules and their absorption cross-sections in vibrational spectroscopy.

A promising route to overcome the limitations of local ML models is to include known physical interactions explicitly [19, 21, 22]. For example, a description of long-range electrostatics can be obtained by

learning atom-centered charge distributions (e.g. atomic charges, dipoles or partitioned electron densities) [19, 23–28]. A prominent recent example of this is the MuML dipole model of Veit *et al*, which uses atomic charges and atom-centered dipoles to predict molecular dipole moments [20]. This idea takes advantage of the fact that the charge distributions around atoms can be predicted with reasonable accuracy from local environments, even if their interactions are long-ranged. While this solves some of the issues of local interatomic potentials, there are also significant downsides: firstly, charge conservation of the overall system is generally not ensured, and secondly non-local charge transfer (e.g. through conjugated π -systems) is not captured [29].

These issues can be fundamentally addressed by switching the target of the ML model: instead of predicting the charge distribution directly, one can predict a charge-dependent energy expression. The charge distribution is then obtained by minimizing this energy expression under the constraint that the charge is conserved. This idea is closely related to classical charge equilibration approaches like QEq [30]. In this manner, charge conservation is rigorously ensured, the description of non-local charge transfer is enabled and a simple route to analytical derivatives is provided through a Hellmann-Feynman-like approach. The advantage of this approach, compared to directly predicting the charge density, can perhaps be understood in analogy to the choice of initial guess in density functional theory (DFT) calculations: while it is common practice to construct the initial guess from a superposition of atomic densities, it has been found that the superposition of atomic potentials yields a significantly improved starting point [31].

So far only few examples of ML-based charge equilibration models have been reported, however. Most notably, Goedecker and co-workers applied a NN-based QEq model to ionic crystals [32, 33]. The corresponding models were trained to predict the energies and forces of reference DFT calculations, using the corresponding partial charges merely as an intermediate quantity. More recently, Behler, Goedecker and co-workers combined this approach with local NN potentials for the description of organic molecules and MgO surfaces [29, 34]. Here, the charge equilibration models were trained on partial charges from reference DFT calculations. Finally, Xie, Persson and Small applied a more flexible charge-dependent NN to describe lithium hydride nanoparticles, using training data from constrained DFT calculations [35].

Herein, we present a new kernel-based approach to charge equilibration termed kernel charge equilibration (kQEq). These models are directly trained on molecular dipole moments and thus avoid the ambiguity associated with choosing population analysis or projection approaches required in other methods. A closed-form linear algebra expression for training kQEq models is derived and their accuracy is benchmarked on the prediction of molecular dipole moments. Finally, limitations and possible extensions are discussed.

2. Theory

2.1. Charge equilibration

Different conventional (i.e. non-ML) charge equilibration and electronegativity equalization methods have been proposed in the literature [30, 36–45]. In the derivation of the charge equilibration approach we largely follow the formalism of Goedecker and coworkers [29, 32–34], which is in turn based on the QEq method of Rappé and Goddard [30, 46]. In this context, QEq can be understood as a kind of semi-empirical, orbital-free DFT, where the electron density $\rho(\mathbf{r})$ is expanded as:

$$\rho(\mathbf{r}) = \rho_0(\mathbf{r}) + \delta\rho(\mathbf{r}), \quad (1)$$

where, $\rho_0(\mathbf{r})$ is a reference density (here the superposition of isolated atom densities) and $\delta\rho(\mathbf{r})$ is a fluctuation term, which describes charge transfer and polarization in the interacting system. We expand $\delta\rho$ into a linear combination of normalized $1s$ Gaussians centered at the atomic positions \mathbf{r}_A and of width α_A

$$\delta\rho(\mathbf{r}) \approx \sum_{A=1}^N -q_A \frac{1}{\pi^{\frac{3}{2}} \alpha_A^3} \exp\left(-\frac{|\mathbf{r} - \mathbf{r}_A|^2}{\alpha_A^2}\right), \quad (2)$$

where N is the number of atoms and q_A are the expansion coefficients. Note that we use the negative of the expansion coefficients q_A here, so that these can directly be interpreted as atomic partial charges. With this approximation, the electron density is completely defined via the charges q_A and it remains to find their optimal values.

To this end, a simple form of the charge-dependent electrostatic energy is assumed:

$$E[\rho] = E_0 + \underbrace{\sum_{A=1}^N (\chi_A q_A + \frac{1}{2} \eta_A q_A^2)}_{\text{Site-Energy}} + \frac{1}{2} \underbrace{\iint \frac{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}d\mathbf{r}'}_{\text{Coulomb-Integral}}. \quad (3)$$

Here, E_0 is a charge-independent reference energy, which we set to zero throughout. The second term (labeled ‘Site-Energy’) is the well-known second-order Taylor expansion of the atomic energy with respect to the charge, with the atomic electronegativity χ_A and the hardness η_A [47]. It provides the energetic contribution incurred by adding or removing electron density from a given atom. The third term (labeled ‘Coulomb-Integral’) is the classical Coulomb potential of the fluctuation density $\delta\rho$. This integral can be computed analytically, using the definition of $\delta\rho$ (see equation (2)):

$$\iint \frac{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}d\mathbf{r}' = \sum_{A=1}^N \left(\frac{2\gamma_{AA}}{\sqrt{\pi}} q_A^2 + \sum_{B=1}^N q_A q_B \frac{\text{erf}(\gamma_{AB} r_{AB})}{r_{AB}} \right), \quad (4)$$

with $\gamma_{AB} = \sqrt{(\alpha_A^2 + \alpha_B^2)}^{-1}$.

This allows rewriting equation (3) to:

$$E[\rho] = E_0 + \sum_{A=1}^N \left[\chi_A q_A + \frac{1}{2} \left(\eta_A + \frac{2\gamma_{AA}}{\sqrt{\pi}} \right) q_A^2 \right] + \frac{1}{2} \sum_{A,B} q_A q_B \frac{\text{erf}(\gamma_{AB} r_{AB})}{r_{AB}}, \quad (5)$$

which makes it explicit that $E[\rho]$ only depends on the charges q_A . We may therefore equivalently use the notation $E(q_1, \dots, q_N)$. Note that this equation has the familiar form of the original QEq formulation, with the slight difference that the hardness parameter in QEq implicitly includes the electrostatic idempotential $\frac{2\gamma_{AA}}{\sqrt{\pi}}$, whereas here η_A only describes the non-classical (e.g. exchange-correlation) contributions to the hardness.

Given the definitions of ρ and $E[\rho]$, we now search for the density that minimizes the energy functional under the constraint that the total number of electrons is conserved. From the definition of $\delta\rho$, it can be seen that this is equivalent to the constraint that $\sum_A q_A = Q_{\text{tot}}$, with the total system charge Q_{tot} . This can be achieved by defining the Lagrangian

$$\mathcal{L}(q_1, \dots, q_N, \lambda) = E(q_1, \dots, q_N) + \lambda \left(\sum_{A=1}^N q_A - Q_{\text{tot}} \right). \quad (6)$$

The constrained minimization of the charges can then be performed by setting up a linear system of equations so that:

$$\frac{\partial \mathcal{L}(q_1, \dots, q_N, \lambda)}{\partial q_A} = \sum_{B=1}^N H_{A,B} q_B + \lambda + \chi_A = 0 \quad (7)$$

$$\frac{\partial \mathcal{L}(q_1, \dots, q_N, \lambda)}{\partial \lambda} = \sum_{A=1}^N q_A - Q_{\text{tot}} = 0 \quad (8)$$

with the elements of the hardness matrix \mathbf{H} defined as:

$$H_{A,B} = \begin{cases} \eta_A + \frac{2\gamma_{AA}}{\sqrt{\pi}} & \text{for } A = B, \\ \frac{\text{erf}(\gamma_{AB} r_{AB})}{r_{AB}} & \text{otherwise.} \end{cases} \quad (9)$$

In matrix notation, this linear system can be formulated as:

$$\underbrace{\begin{pmatrix} H_{1,1} & H_{1,2} & \cdots & H_{1,N} & 1 \\ H_{2,1} & H_{2,2} & \cdots & H_{2,N} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ H_{N,1} & H_{N,2} & \cdots & H_{N,N} & 1 \\ 1 & 1 & \cdots & 1 & 0 \end{pmatrix}}_{\bar{\mathbf{H}}} \cdot \underbrace{\begin{pmatrix} q_1 \\ q_2 \\ \vdots \\ q_N \\ \lambda \end{pmatrix}}_{\bar{\mathbf{q}}} = - \underbrace{\begin{pmatrix} \chi_1 \\ \chi_2 \\ \vdots \\ \chi_N \\ -Q_{\text{tot}} \end{pmatrix}}_{\bar{\mathbf{c}}} \quad (10)$$

with the hardness matrix $\bar{\mathbf{H}}$, the charge vector $\bar{\mathbf{q}}$ and the electronegativity vector $\bar{\mathbf{c}}$. Here, the bar-notation is used to indicate that these arrays are extended by one dimension due to the Lagrange multiplier. The corresponding N -dimensional sub-arrays are indicated by \mathbf{H} , \mathbf{q} and \mathbf{c} in the following.

2.2. Kernel charge equilibration

The QEq approach as outlined above only requires three parameters, namely the electronegativity (χ_A), the non-classical contribution to the hardness (η_A) and the atomic size (α_A) for each species in the system. As a flipside of this elegant simplicity, the accuracy and transferability of the QEq method is limited, however. In the kQEq method proposed herein, we follow the basic idea of Goedecker and coworkers to overcome this limitation [32]. This is achieved by allowing the electronegativity of an atom to change as a function of its chemical environment. Importantly, taking advantage of the fact that both QEq and kernel ridge regression (KRR) are formulated as linear problems, we obtain a closed-form expression for training these models.

The environment-dependent electronegativity in kQEq is defined via a kernel regression Ansatz as:

$$\chi_A(\mathbf{p}_A) = \sum_{B=1}^{N_{\text{train}}} w_B k(\mathbf{p}_A, \mathbf{p}_B), \quad (11)$$

where k is a kernel function, \mathbf{p}_A is a representation vector that encodes the chemical environment of atom A , w_B is a regression coefficient and N_{train} is the number of atoms in the training set. We use the SOAP kernel and representation vector, which are widely used in the construction of interatomic potentials and as descriptors of local environments [18]. We refer to the original literature and a recent review for a detailed account of the corresponding theory and implementation [1, 11, 48]. In general, the kernel function measures the similarity of chemical environments and is defined as:

$$k(\mathbf{p}, \mathbf{p}') = (\mathbf{p} \cdot \mathbf{p}')^2. \quad (12)$$

To derive an expression for the regression coefficients w , we begin by noting that the prediction of charges with QEq can be expressed as a matrix multiplication of the matrix \mathbf{A} with the vector of electronegativities \mathbf{c} :

$$\mathbf{q} = -\mathbf{A}\mathbf{c} \quad (13)$$

where \mathbf{A} is the N -dimensional submatrix of $\bar{\mathbf{H}}^{-1}$. Using equation (11), \mathbf{c} can also be written in terms of a matrix-vector multiplication:

$$\mathbf{c} = \mathbf{K}\mathbf{w} \quad (14)$$

so that

$$\mathbf{q} = -\mathbf{A}\mathbf{K}\mathbf{w}. \quad (15)$$

Here, \mathbf{K} is the kernel matrix (quantifying the similarity between the atoms in the system of interest and the atoms in the training set) and \mathbf{w} is the vector of regression coefficients. There are in principle several options for defining the 'optimal' regression coefficients. One could, e.g. fit them to partial charges obtained from some partitioning of the DFT density. However, the choice of a partial charge model is necessarily somewhat arbitrary and does not guarantee an accurate description of electrostatic interactions. We therefore instead use molecular dipole moments $\boldsymbol{\mu}$ as a reference, which are unambiguously defined physical observables for finite systems and define the leading order term of molecular interactions in the long-range limit.

In a first-principles calculation, the dipole moment is calculated as:

$$\boldsymbol{\mu} = \int \mathbf{r}n(\mathbf{r})d\mathbf{r}, \quad (16)$$

where $n(\mathbf{r})$ is the total charge density (obtained from the sum of the electron and nuclear charge distributions). For a charge equilibration model like kQEq, this simplifies to

$$\boldsymbol{\mu} = \sum_{A=1}^N \mathbf{r}_A q_A = \mathbf{R}\mathbf{q}, \quad (17)$$

using the $3 \times N$ matrix \mathbf{R} with columns \mathbf{r}_A . Note that for an unambiguous definition of $\boldsymbol{\mu}$ for charged systems, center-of-mass shifted coordinates are used throughout.

Combining equations (17) and (15), we obtain:

$$\boldsymbol{\mu} = -\mathbf{R}(\mathbf{A}\mathbf{K}\mathbf{w}), \quad (18)$$

for the kQEq dipole moment. To determine the regression weights w , we then set up a regularized least-squares problem with the loss-function:

$$L = \|(\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{ref}})\|^2 + \sigma \mathbf{w}^T \mathbf{K} \mathbf{w} = \| -\mathbf{R}(\mathbf{A} \mathbf{K} \mathbf{w}) - \boldsymbol{\mu}_{\text{ref}} \|^2 + \sigma \mathbf{w}^T \mathbf{K} \mathbf{w}, \quad (19)$$

where σ is a regularization hyperparameter and the term $\mathbf{w}^T \mathbf{K} \mathbf{w}$ comes from the use of Tikhonov regularization in a kernel regression model. Note that we use the simplest form of regularization, with a single parameter σ to control for overfitting. In principle different regularization strengths could be used for each training point (e.g. proportional to the dipole magnitude).

In a final step, we set $\frac{dL}{dw} = 0$ to minimize L and solve for \mathbf{w} , to obtain:

$$\mathbf{w} = -(\mathbf{A}^T \mathbf{R}^T \mathbf{R} \mathbf{A} \mathbf{K} + \sigma \mathbb{1})^{-1} \mathbf{A}^T \mathbf{R}^T \boldsymbol{\mu}_{\text{ref}}. \quad (20)$$

The above equations are formulated for a single kQEq problem (i.e. a single molecule or system). In practice we train on multiple systems at once. This can still be achieved with a single linear algebra equation by using blocked matrices for \mathbf{A} and \mathbf{R} , and by concatenating the dipole vector elements of all training systems into a single vector.

It should be noted that dipole moments in principle do not contain sufficient information to obtain a unique set of atomic partial charges. One advantage of the kQEq framework is therefore that it offers a natural way to enforce physical constraints on the partial charges. These constraints come in two forms, namely that atoms with similar chemical environments must display similar electronegativities (enforced via the Kernel and regularization) and that the charges must minimize the kQEq energy expression and sum to the total charge (enforced via the QEq framework). Importantly, the set of charges that minimizes a given kQEq energy expression is unique.

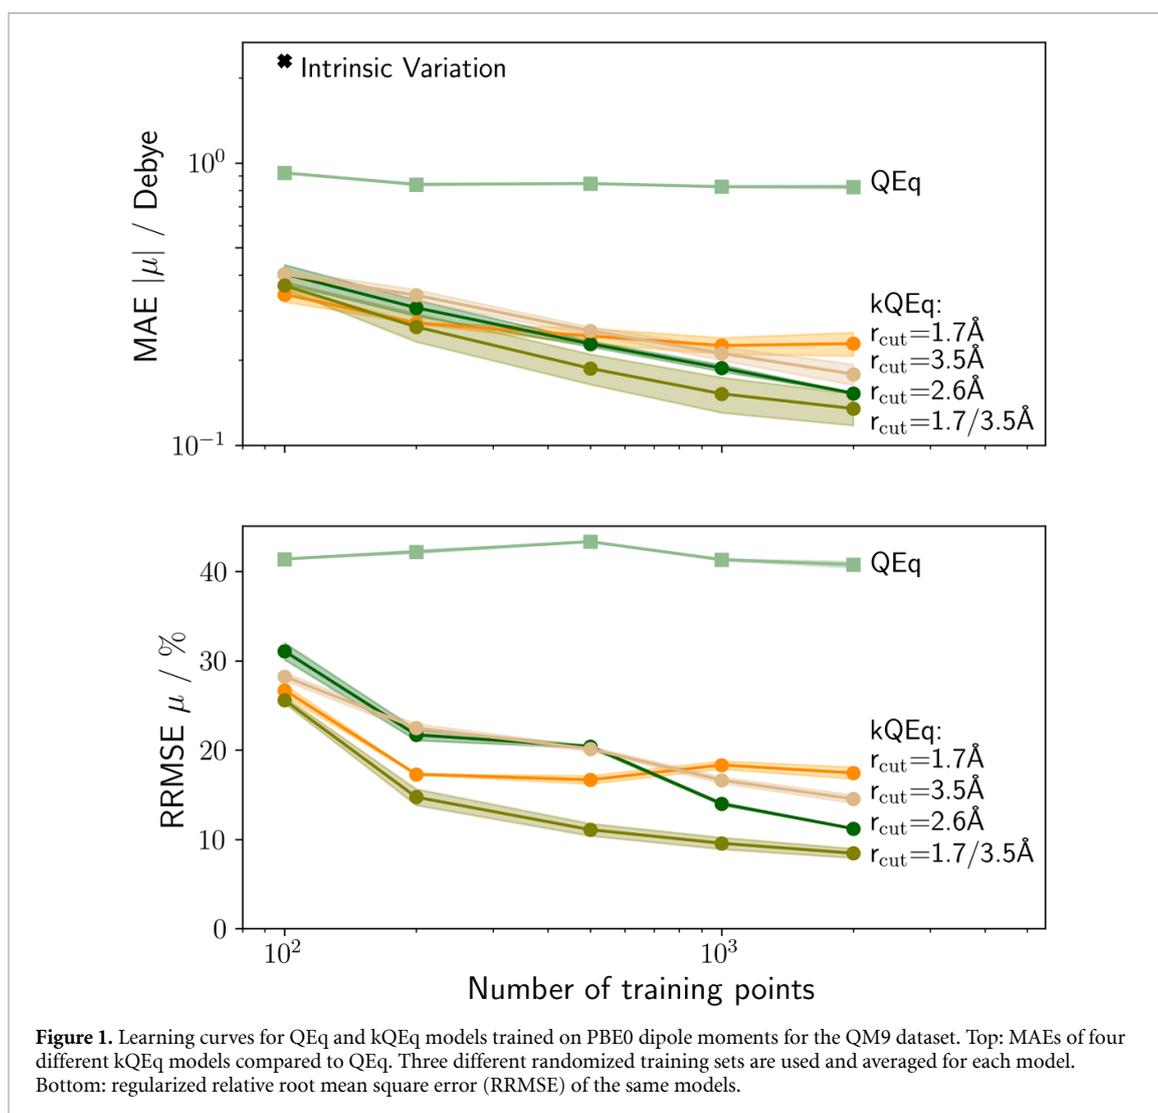
2.3. Hyperparameters

Up to now, an environment-dependent description of the atomic electronegativity χ_A is defined, which can be learned from data. It remains to specify the non-classical contribution to the atomic hardness η_A and the atomic radius α_A for each element. Herein, we choose these by very simple heuristics: α_A is set to be proportional to the original QEq radius of the element in question. These radii are tabulated for all elements up to Lawrencium ($Z = 103$) [30]. Empirically, we found that scaling these radii with single global scaling parameter $s_{\text{at}} = 0.75$ yields satisfactory results. Similarly, the non-classical hardness parameter η_A is set to zero throughout, as we found that this yields robust models while keeping the empiricism of the method as low as possible. These choices are quite simplistic and further optimization would certainly be possible. As shown below, already these simple defaults provide highly accurate results for the investigated systems though.

The main hyperparameters to be considered for SOAP are the cutoff radius r_{cut} within which the neighborhood is expanded and the broadness of the Gaussians used to smear out the atomic positions (σ_{atom}). The choice of these lengthscales governs the range in which the environment of an atom affects its electronegativity and how sensitive it is to changes of the atomic positions. In the following, we keep the ratio between these parameters constant ($\sigma_{\text{atom}} = \frac{r_{\text{cut}}}{8}$) so that for larger cutoffs, the atomic positions are smeared out more strongly. The idea of keeping this ratio fixed is based on the fact that the expressiveness of a given atom-centered basis set is limited by the number of basis functions, meaning that it can either provide a high-resolution picture that is short-ranged or a lower-resolution picture that is longer-ranged. Alternatively, one could increase the number of basis functions for larger cutoff radii instead, but this would lead to a significantly increased computational cost. The particular constant of proportionality we use was found to work well empirically. Note that the `Dscribe` implementation of SOAP is used [48]. Full details are provided in the supplementary information (available online at stacks.iop.org/MLST/3/015032/mmedia). The influence of the cutoff parameter is discussed below.

2.4. Error metrics

To quantify the performance of the kQEq models for predicting molecular dipole moments, we use two complementary metrics. On one hand, we use the mean absolute error (MAE) of predicted absolute dipole moments. This is a common measure of accuracy, which allows direct comparison with previous models. Additionally, we use the root mean squared regularized relative error (RRMSE) as used by Hait and Head-Gordon in [49]. This metric is defined as $\frac{|\mu - \mu_{\text{ref}}|}{\max(\mu_{\text{ref}}, 1 \text{ D})} \times 100\%$, with an arbitrary threshold of 1 Debye that discriminates between small and large dipole moments. In this way, a seamless transition from absolute error (for small dipoles) to relative error (for large dipoles) is achieved, which is necessary since the pure relative error is otherwise severely distorted toward systems with small dipoles.



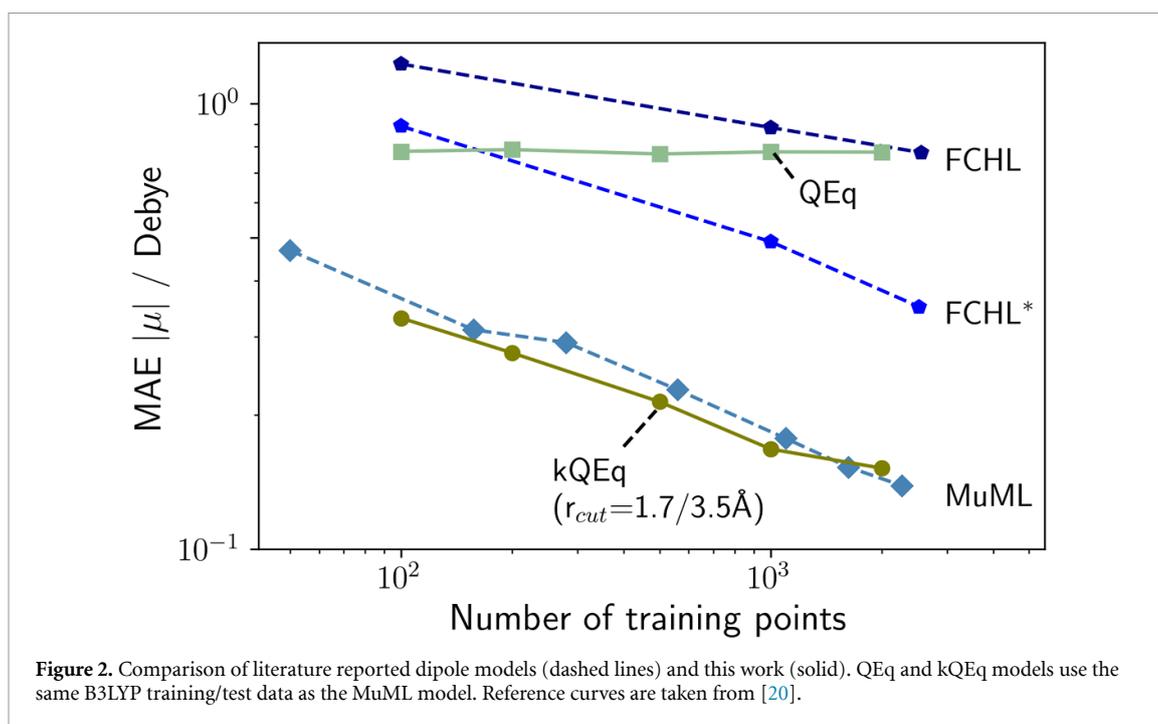
3. Results

3.1. Molecular dipole moments

As a first benchmark we trained kQEq models for predicting dipole moments of organic molecules. As reference data, the dipole moments of 7500 random molecules from the QM9 database were calculated at the PBE0/def2-TZVP level, using ORCA (data provided in the SI) [50–52]. This set spans a wide range of small to medium sized molecules containing the elements C, H, N, O and F. From these structures, we randomly selected a validation set (used to optimize the regularization parameter σ) and a test set of 1000 molecules each. The training sets used below were drawn from the remaining 4000 molecules. Figure 1 depicts the learning curves of different kQEq models, using a range of SOAP cutoffs. For comparison, we also fitted a conventional QEq model to the same data.

All kQEq models clearly outperform the conventional QEq approach, underscoring the benefit of the additional flexibility obtained by using environment-dependent electronegativities. Furthermore, it can be seen that the kQEq models improve continuously when given more data, whereas the MAE of the conventional approach quickly saturates. The model with the smallest SOAP cutoff used here (1.7 Å) shows the best performance for small training sets but stops improving when training on larger sets. Meanwhile, the larger cutoffs we tested (2.6 and 3.5 Å) continuously improve and reach an excellent accuracy of 0.15 D (compared to an intrinsic standard deviation of ca. 2 D).

Following the experience with interatomic potentials based on SOAP, we further tested a 1.7/3.5 Å ‘double SOAP’ representation [7, 53]. This combines a short-ranged/high-resolution kernel with a longer-range/low-resolution one. The corresponding model reaches an even better accuracy and displays the most robust performance across different training set sizes. The latter is particularly evident when considering the RRMSE, which shows that this model consistently improves the relative error when



increasing the training set, whereas some of the other models improve the performance on total dipole moments at the expense of the relative error (i.e. by describing small dipole moments less accurately).

Overall, these results show that the physical description of long-range contributions in kQEq allows the use of rather small cutoffs for the ML part, effectively focusing on the nearest neighbors. This is beneficial both in terms of transferability of the models and the cost of computing the representation. Notably, the philosophically similar BpopNN model of Xie, Persson and Smalls uses a much larger cutoff radius of 13.2 Å [35]. Meanwhile, the MuML dipole model of Veit *et al* (which is also based on SOAP) uses a cutoff of 5 Å [20]. This indicates that kQEq does a good job of partitioning contributions into long-ranged physical terms and a short-ranged ML model.

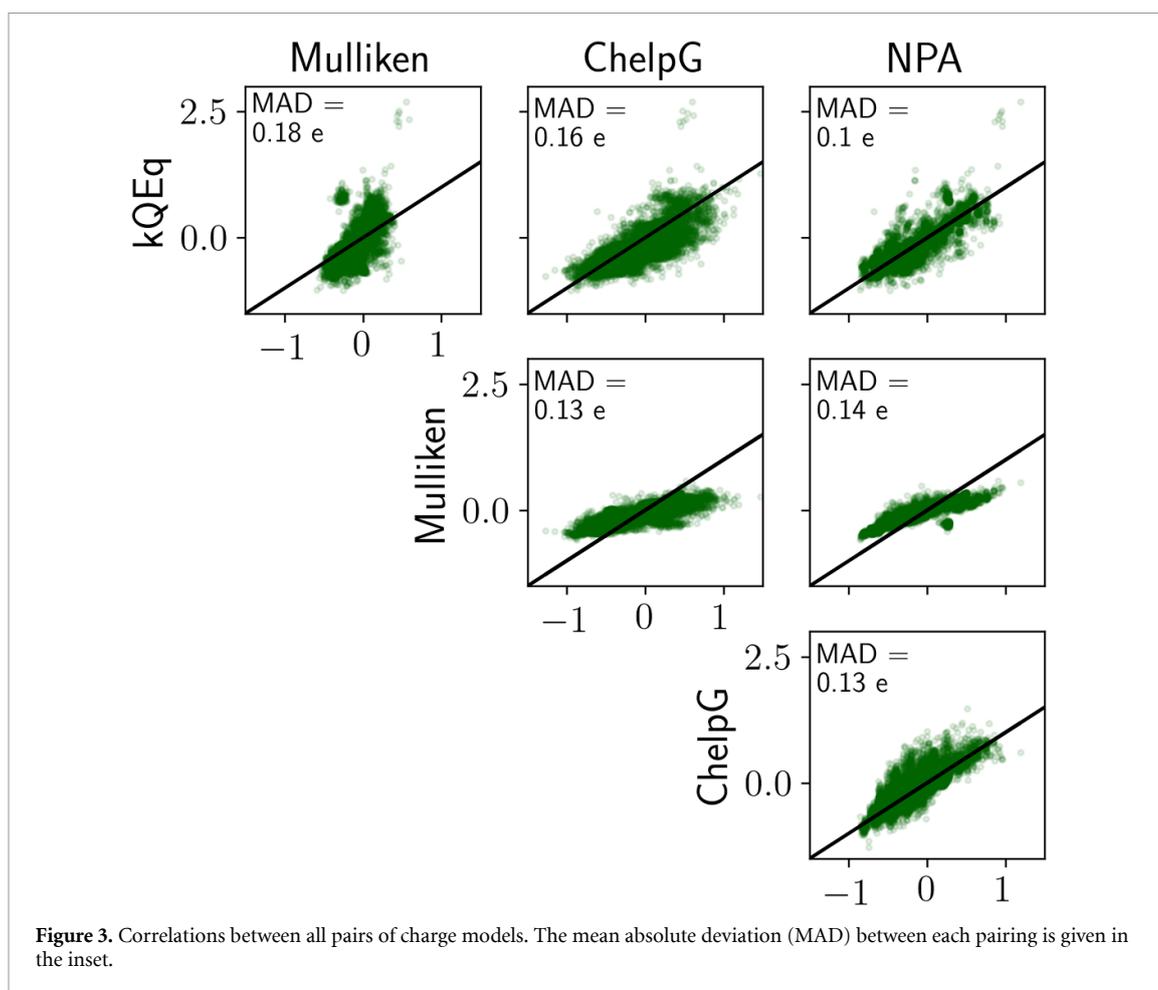
To put this performance into perspective, we compare these models to two recent kernel ML models that are specifically tailored to predicting dipole moments, namely the operator ML approach of Christensen, Faber and Lilienfeld and the aforementioned MuML model of Veit *et al* (see figure 2) [20, 54]. The former uses a modified variant of the Faber-Christensen-Huang-Lilienfeld (FCHL) representation (FCHL*), which can incorporate the response of the ML model to applied electric fields and thus provides a physically rigorous and equivariant route to predicting dipole moments. Meanwhile, the latter uses a decomposition of the total molecular dipole into atomic monopole and dipole contributions, using the equivariant λ -SOAP approach [55]. For reference we also include the learning curve of a naive FCHL model, which simply predicts the total dipole moment as a scalar (taken from [54]).

As already discussed in [54], the FCHL* model is a significant improvement over the scalar approach. It also significantly outperforms conventional QEq for all but the smallest training sets. Meanwhile the MuML and kQEq models display remarkably similar learning curves and represent a further improvement over FCHL*. Note that for consistency the QEq and kQEq models in figure 1 were trained on the same reference data (calculated at the B3LYP level) as the MuML model from [20].

It is also instructive to consider the accuracy of the reference methods themselves. As mentioned above, Hait and Head-Gordon used the RRMSE to benchmark density functional methods against high-level Coupled Cluster [49]. According to this metric, the best kQEq model reaches an accuracy of 8.1% on the dipole moments of QM9. This can be compared with the reported errors of popular hybrid functionals like PBE0 (5.2%) and B3LYP (7%). However, it should be noted that the benchmark in [49] focuses on very small molecules and includes spin-polarized systems, so that this comparison should not be overinterpreted. Nonetheless, it indicates that kQEq models approach hybrid DFT accuracy at a much reduced cost.

3.2. Charge analysis

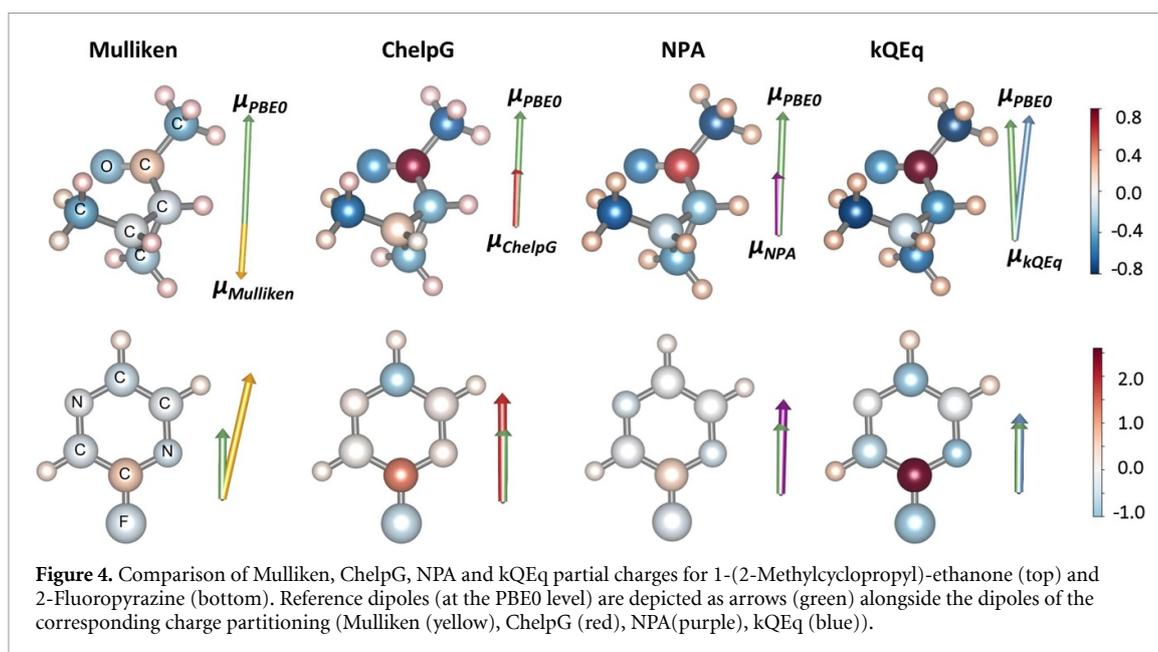
Having established the high accuracy of the kQEq predicted dipole moments, we next turn to the predicted charges themselves. It is well known that the electron density cannot be unambiguously partitioned into atomic partial charges. Consequently, there is no way to objectively measure the quality of such partitionings.



Indeed, there is a fundamental tension between describing the local charge density around each atom versus global properties such as dipole moments or electrostatic potentials, when approximating a continuous charge density with atom-centered spherically symmetric charges. Nonetheless, it is worth considering whether the predicted charges are reasonably intuitive and how they compare to standard population analysis schemes like the one of Mulliken [56], restricted electrostatic potential fits like the ChelpG scheme [57], or localized molecular orbital approaches like the Natural Population Analysis [58, 59]. Here, a double SOAP model trained on 2000 training configurations (at the PBE0 level) is used to provide the kQEq charges (see figure 1).

Figure 3 depicts the pairwise correlations between these partial charge models. As has been noted previously, different charge models in general only display moderate agreement with each other [27]. In fact, the lowest mean absolute deviation (0.1 elementary charges) is found between NPA and kQEq, the largest between Mulliken and kQEq. The latter can be attributed to the fact that kQEq charges appear in a broader range (−1–2.5) while Mulliken charges lie between −0.5 and 0.5. Furthermore, kQEq charges display a much more pronounced clustering into functional groups. This is particularly evident for oxygen. While the Mulliken analysis assigns similar charges to ether and carbonyl oxygen atoms, kQEq predicts the carbonyl groups to be significantly more polar. Similarly, the polarity of nitrile and fluoride functional groups is much higher in kQEq (see SI for element-wise correlation plots with all charge models). In contrast, ChelpG and NPA charges mostly fall into the same range as the kQEq ones. In particular, the large differences for polar functional groups (e.g nitrile or carbonyl functional groups) are not observed here. Nonetheless, large deviations can be seen in other cases, particularly for Fluoride functional groups. In general, the relatively good agreement with NPA charges is in line with the observations of Bultinck and coworkers for conventional charge equilibration methods [44].

Two illustrative examples of these differences are shown for 1-(2-Methylcyclopropyl)-ethanone and 2-Fluoropyrazine in figure 4 (both of which are not part of the kQEq training set). At first glance, the partitionings are qualitatively similar, meaning that the signs of most charges match. However, kQEq, NPA, and ChelpG predict significantly larger absolute charges than Mulliken. In the case of 1-(2-Methylcyclopropyl)-ethanone, this is particularly evident for the carbonyl group, which is much more



strongly polarized. Importantly, these differences have strong consequences for the overall description of the molecular electrostatics. Indeed, the dipole moment calculated from the Mulliken charges points in the opposite direction of the actual dipole. In contrast, the ChelpG and NPA dipoles are well aligned with the reference, but too small in magnitude. Finally, the kQEq charges provide an accurate prediction of both the absolute dipole moment and its direction.

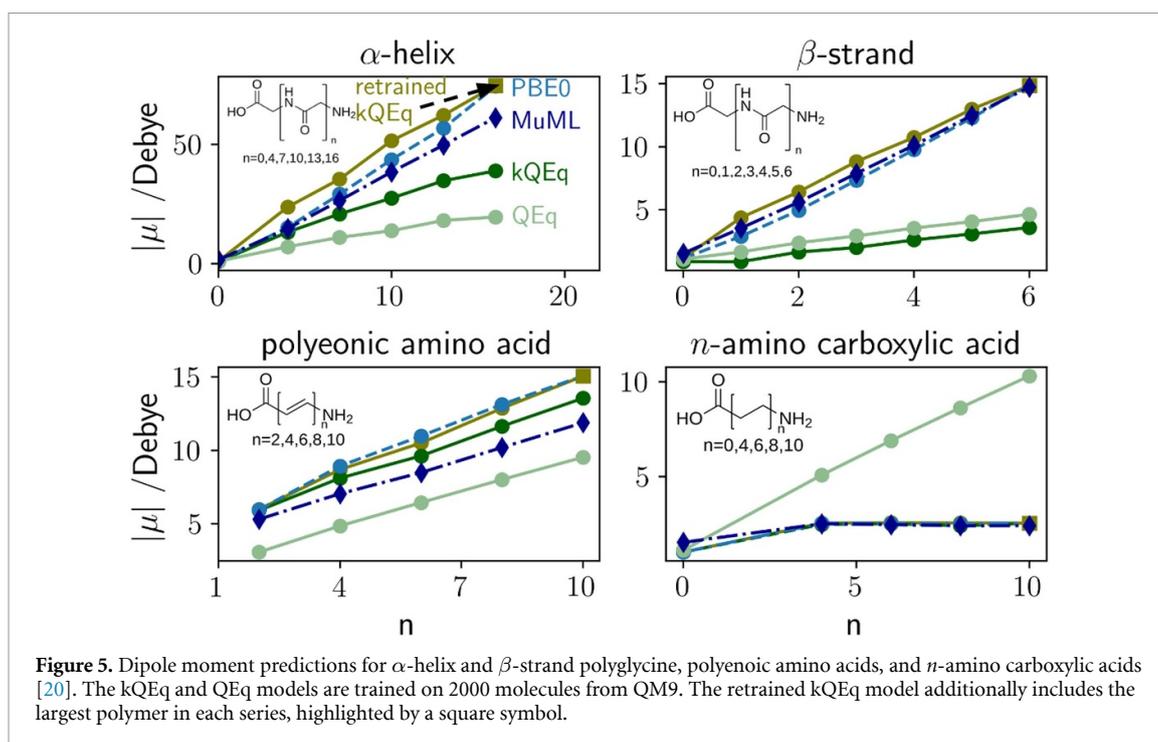
The case of 2-Fluoropyrazine is particularly interesting because it illustrates the observed discrepancy between the charges assigned to C-F groups. Here, kQEq predicts the largest absolute charges on the corresponding C and F atoms. Interestingly, the Mulliken, NPA, and ChelpG charges nonetheless overestimate the molecular dipole moment, while the kQEq prediction provides an excellent fit. This is because the large positive charge on the carbon atom stems not only from charge transfer to Fluorine, but also to the adjacent Carbon and Nitrogen atoms, which partially counterbalance the polarity of the C-F group.

The generally poor correlation between the charges obtained with different schemes raises some questions about the different roles and interpretations partial charges can have. On one hand, they can reflect a local partitioning of the electron density, as in the case of Mulliken, Hirshfeld or Bader charges. On the other hand, they can reflect the electrostatic potential on the surface of a molecules, as in the case of ChelpG and related schemes. While the latter is arguably less arbitrary (as it is directly tied to a physical observable) it has well-known issues with assigning meaningful charges to atoms that are ‘buried’ in bulky molecules.

The kQEq model proposed herein does not neatly fit into these categories. Firstly, it is not a charge partitioning scheme but an ML model. Mulliken, NPA, and ChelpG charges can only be obtained by running a full DFT calculation, whereas the kQEq prediction is much cheaper. Secondly, while kQEq models are trained to reproduce molecular dipole moments, the charges themselves are obtained by minimizing the electrostatic energy expression in equation (5). The Coulomb interactions between partial charges thus provide an important physical constraint on how the charge distribution is approximated. As a consequence, the kQEq derived molecular electrostatic potentials are in good agreement with the ChelpG and NPA ones (see SI). Nonetheless, the kQEq charges are comparatively large. Since bio-organic forcefields usually rely on ChelpG or scaled Mulliken charges, the current kQEq models may thus not be directly applicable in MD simulations. Indeed, the development of interatomic potentials based on kQEq will likely require more complex loss functions, e.g. incorporating energetic contributions or higher order moments. This will be the subject of future work.

3.3. Extrapolating beyond small molecules

A well-known drawback of the conventional QEq approach is that it suffers from a delocalization error akin to that observed for local density functionals. This is e.g. evident in the fact that QEq models incorrectly dissociate molecules into partially charged atoms, since electronegativity differences between isolated atoms lead to spurious charge transfer [39]. While kQEq could in principle cure this particular pathology (by assigning the same electronegativity to all isolated atoms), the more general delocalization tendencies of QEq will likely be inherited by kQEq to an extent. In this section we explore this issue by testing the performance of kQEq for the organic polymer chains dataset introduced by Veit *et al* [20] (see figure 5).



This dataset consists of two types of structures. On one hand, glycine polypeptides in the α -helix and β -strand configurations are considered. On the other hand, polyenoic amino acids and n -amino carboxylic acids are included, which consist of a carboxylic acid and an amine group separated by a conjugated double bond or alkane spacer, respectively. Each of these systems shows characteristic changes of the total dipole moment as the polymer length increases. For the polypeptides, each additional amide bond is itself polar, so that the total dipole increases approximately linearly with the system size. However, the precise behavior depends on the spatial orientation of these bond dipoles and their interactions, so that the α -helix and β -strand configurations show different scalings. The polyenoic amino acid chains also display a linearly increasing dipole moment. In this case this is not due to the addition of polar bonds, however, but due to the polarization of the delocalized π -electrons in the spacer. Finally, the dipole of the n -amino carboxylic acids remains constant upon increasing the chain length, since no polar bonds or delocalized electrons are present.

We again use a ‘double SOAP’ kQEq model trained on 2000 randomly drawn QM9 molecules for comparison. This model has a mixed performance for this test. For polyenoic and n -amino carboxylic acids, the performance is very good. This is both in terms of the qualitative features (linear increase in dipole moments, vs. quick saturation) and the quantitative agreement. Furthermore, kQEq is a strong improvement over conventional QEq here. In contrast, the performance for the polypeptides is less satisfactory. While the linear trend is correctly captured, the magnitude of the dipoles is significantly underestimated, in particular for the β -strand. This behavior is analogous to the underestimation of dipoles typically observed with local functionals, due to the delocalization error [49]. While this shows that kQEq does not automatically cure all pathologies of QEq, a clear advantage of a ML approach is that it can be improved with more data. Indeed, by including the longest chains of each type explicitly in the training set, a retrained kQEq model can be generated that captures these trends more accurately, see figure 5. Potentially, an improved extrapolation behavior could be obtained by using a more flexible expression for the site-energy in equation (3), effectively tackling the problem at its root.

Figure 5 also shows the best MuML model from [20] for comparison. This model performs quite well overall, in particular for the polypeptide systems, with slightly worse performance for the carboxylic acids. In this context, it is worth discussing the fundamental differences between MuML and kQEq. The former uses a fully local decomposition of the overall dipole moment. This works very well for situations where the charge distribution of a large system is essentially just a sum of smaller fragments (as for the polypeptides). In contrast, long-range charge transfer and polarization effects cannot be described by such a model. Specifically, in a system where two functional groups A and B are separated by some spacer, the MuML charge on A will be independent of B . This is reflected in the underestimated dipole moments of longer polyenoic amino acid chains. QEq and kQEq are in principle able to describe such non-local effects. The QEq curve for n -amino carboxylic acid reveals that this is not necessarily an advantage however, as large

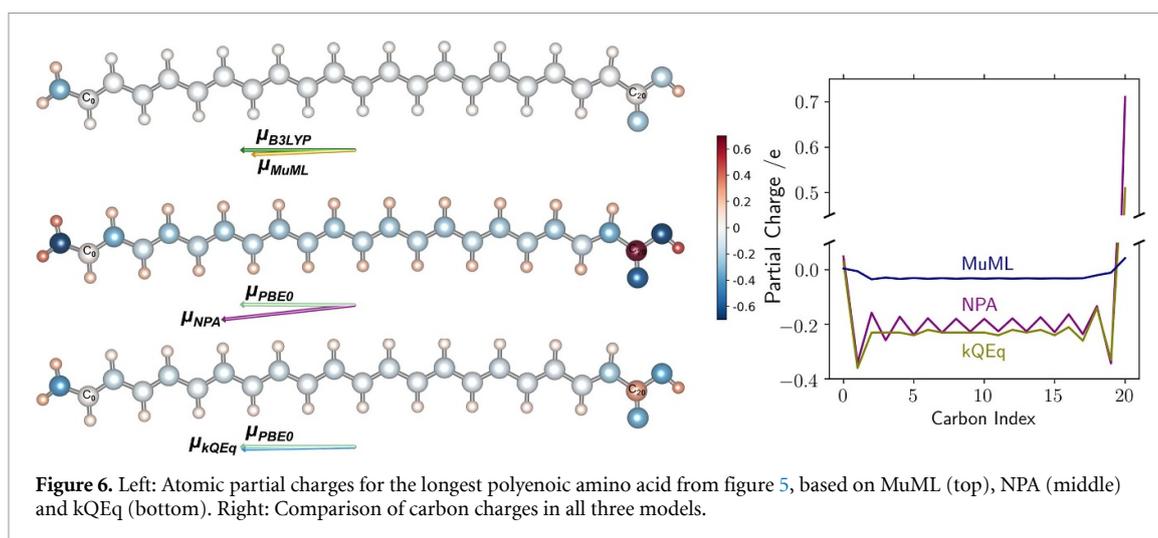


Figure 6. Left: Atomic partial charges for the longest polyenoic amino acid from figure 5, based on MuML (top), NPA (middle) and kQEq (bottom). Right: Comparison of carbon charges in all three models.

unphysical charge transfer is predicted in this case. In other words, a purely local model will generally lead to more systematic and less dramatic failures than a poor non-local one. Fortunately, kQEq and related methods now provide a framework for sophisticated non-local charge equilibration models.

More detailed insight into the charge distributions of the different models can be obtained by comparing the partial charges directly (see figure 6). Here, the MuML and kQEq charges for the longest polyenoic amino acid are shown. For comparison, quantum mechanical NPA charges are also included. This reveals that the kQEq charges display a much better qualitative agreement with NPA. In particular, the oscillatory behavior of the carbon charges (which can be attributed to polarization effects) is completely absent in MuML but clearly observable for kQEq. This is notable, given that both ML models are exclusively trained on dipole moments (i.e. without any atomistic detail) and supports the notion that the kQEq framework introduces useful physical constraints on the charges.

More generally, it should be noted that kQEq is not primarily intended as a stand-alone molecular dipole model. Since it is based on an energy functional, it can be used to model long-range electrostatic interactions in combination with local interatomic potentials. Indeed, it remains an open question whether dipole moments alone provide enough information for this purpose. Fortunately, the current approach can easily be expanded to also include higher moments, electrostatic potentials or reference partial charges. This will be explored in future work.

4. Conclusion

In this work, we introduced kQEq, a kernel-based approach to charge equilibration in molecules. In contrast to conventional charge equilibration methods like QEq, a data-driven, environment-dependent description of atomic electronegativities is introduced. kQEq models trained on molecular dipole moments display excellent performance, *en par* with or better than state-of-the-art kernel models, specifically tuned to predicting dipole moments [20, 54].

The formalism presented herein opens the door toward physics-based kernel ML models for predicting atomic charges, to be used in combination with reactive interatomic potentials. Importantly, the presented approach is quite general and can be extended to other fit targets (e.g. quadrupole moments and electrostatic potentials) and to more flexible density representations (e.g. using atom centered dipoles in addition to partial charges).

While this work focuses on molecular systems, the application to inorganic materials is also envisioned. For finite systems (e.g. nanoparticles) this is in principle straightforward. The use of other fit targets may become important for larger systems, however, as the dipole moment alone might contain too little information in this case. For periodic systems multipole moments are in general ill-defined. Here, the current approach will have to be extended to reproduce other electronic properties.

It should also be noted that QEq-based frameworks are likely not equally well suited for different kinds of materials. In principle, one would expect the best performance for metallic systems, where charges are strongly delocalized and mobile. In contrast, polar insulators or interfaces may be less well described due to the intrinsic delocalization error of charge equilibration models. More general site-energy expressions could be developed to overcome these tendencies. Ideally, such developments will ultimately converge with recent

developments in ML-based DFT, yielding a new generation of orbital-free density functionals [60–63]. Work in this direction is ongoing in our groups.

Data availability statement

All data and a reference implementation of the kQEq method is available at <https://gitlab.com/jmargraf/kqeq>.

Acknowledgments

The authors thank the DFG for financial support through Germany's Excellence Strategy—EXC 2089/1-390776260. S W and K R further acknowledge funding from Deutsche Forschungsgemeinschaft (DFG) under project RE1509/18-2.

ORCID iD

Johannes T Margraf  <https://orcid.org/0000-0002-0862-5289>

References

- [1] Deringer V L *et al* 2021 Gaussian process regression for materials and molecules *Chem. Rev.* **121** 10073–141
- [2] Behler J 2021 Four generations of high-dimensional neural network potentials *Chem. Rev.* **121** 10037–72
- [3] Unke O T *et al* 2021 Machine learning force fields *Chem. Rev.* **121** 10142–86
- [4] Manzhos S and Carrington T J 2020 Neural network potential energy surfaces for small molecules and reactions *Chem. Rev.* **121** 10187–10217
- [5] Rupp M, Tkatchenko A, Müller K-R and von Lilienfeld O A 2012 Fast and accurate modeling of molecular atomization energies with machine learning *Phys. Rev. Lett.* **108** 058301
- [6] Deringer V L, Caro M A and Csányi G 2019 Machine learning interatomic potentials as emerging tools for materials science *Adv. Mater.* **31** 1902765
- [7] Bartók A P *et al* 2017 Machine learning unifies the modeling of materials and molecules *Sci. Adv.* **3** e1701816
- [8] Bartók A P, Payne M C, Kondor R and Csányi G 2010 Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons *Phys. Rev. Lett.* **104** 136403
- [9] Behler J and Parrinello M 2007 Generalized neural-network representation of high-dimensional potential-energy surfaces *Phys. Rev. Lett.* **98** 146401
- [10] Artrith N, Morawietz T and Behler J 2011 High-dimensional neural-network potentials for multicomponent systems: applications to zinc oxide *Phys. Rev. B* **83** 153101
- [11] Bartók A P, Kondor R and Csányi G 2013 On representing chemical environments *Phys. Rev. B* **87** 184115
- [12] Faber F A, Christensen A S, Huang B and von Lilienfeld O A 2018 Alchemical and structural distribution based representation for universal quantum machine learning *J. Chem. Phys.* **148** 241717
- [13] Christensen A S, Bratholm L A, Faber F A and Anatolev O 2020 FCHL revisited: faster and more accurate quantum machine learning *J. Chem. Phys.* **152** 044107
- [14] Schütt K T, Arbabzadah F, Chmiela S, Müller K R and Tkatchenko A 2017 Quantum-chemical insights from deep tensor neural networks *Nat. Commun.* **8** 13890
- [15] Hy T S, Trivedi S, Pan H, Anderson B M and Kondor R 2018 Predicting molecular properties with covariant compositional networks *J. Chem. Phys.* **148** 241745
- [16] Klicpera J, Groß J and Günnemann S 2020 Directional message passing for molecular graphs (arXiv:200303123 [Phys. Stat])
- [17] Klicpera J, Giri S, Margraf J T and Günnemann S 2020 Fast and uncertainty-aware directional message passing for non-equilibrium molecules (arXiv:201114115 [physics])
- [18] Cheng B *et al* 2020 Mapping materials and molecules *Acc. Chem. Res.* **53** 1981–91
- [19] Unke O T and Meuwly M 2019 PhysNet: a neural network for predicting energies, forces, dipole moments and partial charges *J. Chem. Theory Comput.* **15** 3678–93
- [20] Veit M, Wilkins D M, Yang Y, DiStasio R A J and Ceriotti M 2020 Predicting molecular dipole moments by combining atomic partial charges and atomic dipoles *J. Chem. Phys.* **153** 024113
- [21] Yao K, Herr J E, Toth D W, McKintyre R and Parkhill J 2018 The TensorMol-0.1 model chemistry: a neural network augmented with long-range physics *Chem. Sci.* **9** 2261–9
- [22] Staacke C G *et al* 2021 On the role of long-range electrostatics in machine-learned interatomic potentials for complex battery materials *ACS Appl. Energy Mater.* **4** 12562–9
- [23] Mahmoud C B, Anelli A, Csányi G and Ceriotti M 2020 Learning the electronic density of states in condensed matter *Phys. Rev. B* **102** 235130
- [24] Fabrizio A, Grisafi A, Meyer B, Ceriotti M and Corminboeuf C 2019 Electron density learning of non-covalent systems *Chem. Sci.* **10** 9424–32
- [25] Zubatyuk R, Smith J S, Leszczynski J and Isayev O 2019 Accurate and transferable multitask prediction of chemical properties with an atoms-in-molecules neural network *Sci. Adv.* **5** 10
- [26] Gastegger M and Marquetand P 2020 Molecular dynamics with neural network potentials *Machine Learning Meets Quantum Physics (Lecture Notes in Physics)* ed K T Schütt *et al* (Cham: Springer Int. Publishing) pp 233–52
- [27] Sifain A E *et al* 2018 Discovering a transferable charge assignment model using machine learning *J. Phys. Chem. Lett.* **9** 4495–501
- [28] Nebgen B *et al* 2018 Transferable dynamic molecular charge assignment using deep neural networks *J. Chem. Theory Comput.* **14** 4687–98
- [29] Ko T W, Finkler J A, Goedecker S and Behler J 2021 A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer *Nat. Commun.* **12** 398

- [30] Rappe A K and Goddard W A 1991 Charge equilibration for molecular dynamics simulations *J. Phys. Chem.* **95** 3358–63
- [31] Lehtola S 2019 Assessment of Initial Guesses for self-consistent field calculations. Superposition of atomic potentials: simple yet efficient *J. Chem. Theory Comput.* **15** 1593–604
- [32] Ghasemi S A, Hofstetter A, Saha S and Goedecker S 2015 Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network *Phys. Rev. B* **92** 045131
- [33] Faraji S et al 2017 High accuracy and transferability of a neural network potential through charge equilibration for calcium fluoride *Phys. Rev. B* **95** 1–11
- [34] Ko T W, Finkler J A, Goedecker S and Behler J 2021 General-purpose machine learning potentials capturing nonlocal charge transfer *Acc. Chem. Res.* **54** 808–17
- [35] Xie X, Persson K A and Small D W 2020 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** 4256–70
- [36] Gasteiger J and Marsili M 1980 Iterative partial equalization of orbital electronegativity—a rapid access to atomic charges *Tetrahedron* **36** 3219–28
- [37] Rick S W, Stuart S J and Berne B J 1994 Dynamical fluctuating charge force fields: application to liquid water *J. Chem. Phys.* **101** 6141–56
- [38] van Duin A C T, Dasgupta S, Lorant F and Goddard W A 2001 ReaxFF: a reactive force field for hydrocarbons *J. Phys. Chem. A* **105** 9396–409
- [39] Morales J and Martínez T J 2001 Classical fluctuating charge theories: the maximum entropy valence bond formalism and relationships to previous models *J. Phys. Chem. A* **105** 2842–50
- [40] Morales J and Martínez T J 2004 A new approach to reactive potentials with fluctuating charges: quadratic valence-bond model *J. Phys. Chem. A* **108** 3076–84
- [41] Nistor R A, Polihronov J G, Müser M H and Mosey N J 2006 A generalization of the charge equilibration method for nonmetallic materials *J. Chem. Phys.* **125** 094108
- [42] Valone S M 2011 Quantum mechanical origins of the Iczkowski-Margrave model of chemical potential *J. Chem. Theory Comput.* **7** 2253–61
- [43] Pracht P, Caldeweyher E, Ehlert S and Grimme S 2019 A robust non-self-consistent tight-binding quantum chemistry method for large molecules *ChemRxiv* 1–19
- [44] Bultinck P et al 2002 The electronegativity equalization method II: applicability of different atomic charge schemes *J. Phys. Chem. A* **106** 7895–901
- [45] Bultinck P et al 2002 The electronegativity equalization method I: parametrization and validation for atomic charge calculations *J. Phys. Chem. A* **106** 7887–94
- [46] Kwon S, Naserifar S, Lee H M and Goddard W A 2018 Polarizable charge equilibration model for transition-metal elements *J. Phys. Chem. A* **122** 9350–8
- [47] Parr R G and Pearson R G 1983 Absolute hardness: companion parameter to absolute electronegativity *J. Am. Chem. Soc.* **105** 7512–16
- [48] Himanen L et al 2019 Dscribe: library of descriptors for machine learning in materials science *Comput. Phys. Commun.* **247** 106949
- [49] Hait D and Head-Gordon M 2018 How accurate is density functional theory at predicting dipole moments? An assessment using a new database of 200 benchmark values *J. Chem. Theory Comput.* **14** 1969–81
- [50] Ramakrishnan R, Dral P O, Rupp M and von Lilienfeld O A 2014 Quantum chemistry structures and properties of 134 kilo molecules *Sci. Data* **1** 1–7
- [51] Neese F 2012 The ORCA program system *Wiley Interdiscip. Rev.-Comput. Mol. Sci.* **2** 73–8
- [52] Neese F 2018 Software update: the ORCA program system, version 4.0 *Wiley Interdiscip. Rev.-Comput. Mol. Sci.* **8** e1327
- [53] Stocker S, Csányi G, Reuter K and Margraf J 2020 Machine learning in chemical reaction space *Nat. Commun.* **11** 5505
- [54] Christensen A S, Faber F A and von Lilienfeld O A 2019 Operators in quantum machine learning: response properties in chemical space *J. Chem. Phys.* **150** 064105
- [55] Grisafi A, Wilkins D M, Csányi G and Ceriotti M 2018 Symmetry-adapted machine learning for tensorial properties of atomistic systems *Phys. Rev. Lett.* **120** 36002
- [56] Mulliken R S 1955 Electronic population analysis on LCAO–MO molecular wave functions. i *J. Chem. Phys.* **23** 1833–40
- [57] Breneman C M and Wiberg K B 1990 Determining atom-centered monopoles from molecular electrostatic potentials. The need for high sampling density in formamide conformational analysis *J. Comput. Chem.* **11** 361–73
- [58] Reed A E, Weinstock R B and Weinhold F 1985 Natural population analysis *J. Chem. Phys.* **83** 735–46
- [59] Nikolaienko T Y, Bulavin L A and Hovorun D M 2014 JANPA: an open source cross-platform implementation of the natural population analysis on the Java platform *Comput. Theor. Chem.* **1050** 15–22
- [60] Brockherde F et al 2017 Bypassing the Kohn-Sham equations with machine learning *Nat. Commun.* **8** 872
- [61] Bogojeski M, Vogt-Maranto L, Tuckerman M E, Müller K-R and Burke K 2020 Quantum chemical accuracy from density functional approximations via machine learning *Nat. Commun.* **11** 5223
- [62] Dick S and Fernandez-Serra M 2020 Machine learning accurate exchange and correlation functionals of the electronic density *Nat. Commun.* **11** 3509
- [63] Margraf J T and Reuter K 2021 Pure non-local machine-learned density functional theory for electron correlation *Nat. Commun.* **12** 344