

Exploration and Validation of Force Field Design Protocols through QM-to-MM Mapping

Chris Ringrose,[†] Joshua T. Horton,[†] Lee-Ping Wang,[‡] and Daniel J. Cole^{*,†}

[†]*School of Natural and Environmental Sciences, Newcastle University, Newcastle upon
Tyne NE1 7RU, United Kingdom*

[‡]*Department of Chemistry, The University of California at Davis, Davis, California 95616,
United States*

E-mail: daniel.cole@ncl.ac.uk

Abstract

The scale of the parameter optimisation problem in traditional molecular mechanics force field construction means that design of a new force field is a long process, and sub-optimal choices made in the early stages can persist for many generations of the force field. We hypothesise that careful use of quantum mechanics to inform molecular mechanics parameter derivation (QM-to-MM mapping) should be used to significantly reduce the number of parameters that require fitting to experiment and increase the pace of force field development. Here, we design a collection of 15 new protocols for small, organic molecule force field design, and test their accuracy against experimental liquid properties. Our best performing model has only seven fitting parameters, yet achieves mean unsigned errors of just 0.031 g/cm³ and 0.69 kcal/mol in liquid densities and heats of vaporisation, compared to experiment. The software required to derive the designed force fields is freely available at <https://github.com/qubekit/QUBEKit>.

Introduction

Classical molecular mechanics force fields are widely used approximations to recover the energy and forces of an atomistic system as a function of its nuclear coordinates. Force fields are an invaluable companion to quantum mechanical modelling in cases where the number of atoms or the time scale to be modelled would be otherwise restricted. Such techniques have seen applications in, for example, the modelling of battery materials,^{1,2} organic light-emitting diodes,^{3,4} crystal structure prediction,⁵ and particularly in computer-aided drug design where they form the basis of molecular dynamics, docking and free energy simulations for calculating protein-ligand binding affinity.^{6,7}

Typically, for modelling organic molecules in biology and chemistry, the force field is a sum of bonded (including harmonic bond-stretching and angle-bending terms, and anharmonic 4-body torsion potentials) and non-bonded (including Coulombic and Lennard-Jones interactions) terms. Small molecule force fields of this form include GAFF,⁸ CGenFF,⁹ OPLS,¹⁰ and the Open Force Field 1.0.0 (‘Parsley’) force field.¹¹ In general, the bonded and charge parameters of the force field may be fit to quantum mechanical potential energy surfaces and electrostatic potentials, respectively. On the other hand, the Lennard-Jones parameters are nearly always fit to experimental pure liquid properties.^{10,12} To be broadly effective, small molecule force fields require accurate parameterisation of each of the above terms for all chemical space. While recent trends have seen parameters fit to larger datasets, complete coverage is challenging. Numerous studies, such as the SAMPL blind challenges¹³ have shown that there is plenty of room for improvement when it comes to predictive molecular modelling with force fields.

In contrast to the empirical molecular mechanics force fields common in biological modelling, interatomic potentials may instead be derived directly from quantum mechanics (QM). For example, intermolecular perturbation theory may be used to separate the full QM interaction energy into physically motivated components,¹⁴ or machine learning based potentials may be trained on QM energies and forces.¹⁵ These methods are attractive because they

remove the empiricism of the force field approach, but they are computationally more expensive, and unless they fully incorporate, for example, many body and quantum nuclear effects they are unlikely to reproduce condensed phase observables with sufficient accuracy.

QM-to-MM mapping reduces size of parameter search space.

Between the empiricism of typical classical force fields and the accurate, yet expensive, *ab initio* derived force fields, increasing attention is being drawn to mapping physically motivated parameters from QM into simple MM functional forms. By deriving bespoke force field parameters for the molecule under study directly from QM, the number of transferable, empirical parameters to be fit is significantly reduced. By retaining common MM functional forms, the potentials may readily be implemented in widely used MM software and used in, for example, free energy calculations. Bespoke intramolecular bond, angle and, particularly, torsion parameters may be readily derived from a small number of QM calculations either by fitting to Hessian matrices¹⁶ or potential energy surface scans.^{17,18} Atom-centred partial charges can be routinely derived from either semi-empirical calculations,^{19,20} or QM electrostatic potential fitting,^{21,22} or atoms-in-molecule electron density partitioning.^{23–26}

Less common, but perhaps most interestingly, is the possibility of using atoms-in-molecule electron density partitioning to derive other components of the non-bonded interaction from QM, in particular dispersion coefficients (C_6 ,^{27–30} C_8 ,^{31,32} ...), atomic polarisabilities,^{32,33} and off-centre charges to model electronic anisotropy.^{28,29} For example, Visscher and Geerke have derived a polarisable force field model, with a higher order dispersion term derived from iterative Hirshfeld atoms-in-molecule analysis, and applied it to small-molecule amino acid analogues.³⁴ Of the 138 nonbonded parameters in their model, 132 are determined from QM, leaving just six to be fit to experiment. Kantonen et al³⁵ employ a mapping of information from minimal basis iterative stockholder (MBIS) atomic electron densities (in particular the decay constants of the electron densities), to derive atom-specific Lennard-Jones parameters. The mapping parameters (two per element) are fit to experimental liquid properties using

the ForceBalance software.³⁶

Our own work has focused on the development of a QUantum mechanical BEspoke (QUBE) force field and associated toolkit (QUBEKit),²⁹ built on QM-to-MM parameter mapping. QUBE bond and angle force field parameters are derived from the QM Hessian matrix of the molecule under study, using the modified Seminario method.³⁷ Atomic partial charges are computed from the density derived electrostatic and chemical (DDEC) partitioned atomic electron densities.^{38,39} The Tkatchenko-Scheffler method²⁷ is used to derive C_6 parameters from the same atomic electron densities, and the repulsive part of the Lennard-Jones potential is derived from atoms-in-molecule atomic radii.²⁸ Once all other parameters are in place, flexible torsion parameters can be fit to QM dihedral scans using interfaces between QUBEKit and external QM and MM software packages.²⁹ A small number of mapping parameters (in this case free atom radii for use in the derivation of Lennard-Jones parameters) is used to ensure accuracy of condensed phase properties,²⁹ and the resulting force fields have also been shown to perform well in the calculation of protein–ligand binding free energies.^{40–42}

Force field design choices can be tested and optimised.

The above mentioned force fields derived from QM-to-MM parameter mappings have a simple functional form, are straightforward to derive from a small number of simple QM calculations, and are bespoke to the system under study. However, despite the advantages, there are still multiple design choices that must be made when constructing a QM-derived force field. These range from the choice of functional form to use for the final force field, to the choice of underlying QM method to compute the electron density, how to partition it between atoms, and how to map the partitioned density to force field parameters. These choices may be regarded as hyperparameters of the model.⁴³ Drawing on lessons from the machine learning field, separate models for each hyperparameter should be systematically trained and tested. However, these hyperparameters are often ‘baked in’ to the design of the force field, and

correcting known problems can be challenging.⁴⁴ Until recently, the manual training of force field parameters for a given set of hyperparameters (for example, Lennard-Jones parameters for a given charge model) would be extremely time-consuming. However, allied with improvements in data collection and retrieval,⁴⁵ advances in automated parameter fitting to experimental liquid properties now make this possible.^{36,46,47} Most notably, the ForceBalance software enables reproducible and automated force field parameterisation against a set of target data that can include either QM or experimental input data. In the context of the current work, ForceBalance was used to tune the QM-to-MM mapping parameters against experimental liquid data in the aforementioned study by Kantonen et al,³⁵ and was also used to rapidly re-fit Lennard-Jones parameters for a range of proposed implicit solvent models (hyperparameters) in the development of the RESP2 method.²²

An automated toolkit for force field design, training and testing.

Until now, the systematic testing of design choices in MM force fields has been rather limited, and it is almost impossible for individual users to undertake this task due to the scale of the parameter optimisation problem. In what follows, we describe our interface between the QUBEKit and ForceBalance software packages. This open source (including all dependencies) software workflow allows users to make a choice of force field hyperparameters, train the model against experimental liquid properties, and unambiguously test the resulting accuracy. We provide improved protocols for deriving the positions and magnitudes of off-site charges (virtual sites) to model anisotropic electron density from the output of atoms-in-molecule electron density partitioning calculations, as well as tools for including them in torsion parameter fitting. Using a train/test split of 15/51 small, organic molecules, we develop 15 force field protocols that aim to test the accuracy of different choices of i) underlying QM method and basis set, ii) atoms-in-molecule electron density partitioning method, iii) implicit solvent model parameters (used to pre-polarise charges for use in the condensed phase), iv) mapping of atomic electron densities to LJ parameters, v) assignment

of Lennard-Jones parameters to polar hydrogen atoms, and vi) use of virtual sites to model anisotropic electron density. We show that a wide range of different choices can be made in the force field design process, some of which do not affect the overall accuracy, and some of which show marked improvements. For our best performing protocol, we provide all of the parameters required for users to derive their own force fields using the QUBEKit software. With these developments comes the potential for future rapid design of next-generation force fields.

Methods

QUBEKit Workflow

QUBEKit is a largely Python 3.6+ based force field derivation toolkit for mac and Linux operating systems. By combining multiple open-source bioinformatics and quantum chemistry packages into a single workflow, it is possible to generate bespoke force fields for molecular dynamics simulations (Figure 1(a)). QUBEKit can be run using only Anaconda or PyPI installable open-source packages, with the option to use Gaussian⁴⁸ for QM and torsion optimisations, and the fortran package Chargemol^{24,25} for atoms-in-molecule analysis.

The general workflow of QUBEKit begins with generating a set of initial parameters for an OPLS-style force field:

$$\begin{aligned}
 U = & \sum_{Bonds} \frac{k_r}{2}(r - r_0)^2 + \sum_{Angles} \frac{k_\theta}{2}(\theta - \theta_0)^2 \\
 & + \sum_{Dihedrals} \left[\frac{V_1}{2}(1 + \cos(\phi)) + \frac{V_2}{2}(1 - \cos(2\phi)) + \frac{V_3}{2}(1 + \cos(3\phi)) + \frac{V_4}{2}(1 - \cos(4\phi)) \right] \quad (1) \\
 & + \sum_{Pairs} \frac{q_i q_j}{r_{ij}} + 4\epsilon_{ij} \left(\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right)
 \end{aligned}$$

where k_r , r_0 , k_θ , θ_0 , V_{1-4} , q_i , ϵ_{ij} and σ_{ij} are force field parameters to be derived. QUBEKit can take as input most industry-standard molecule formats, such as a SMILES string, or

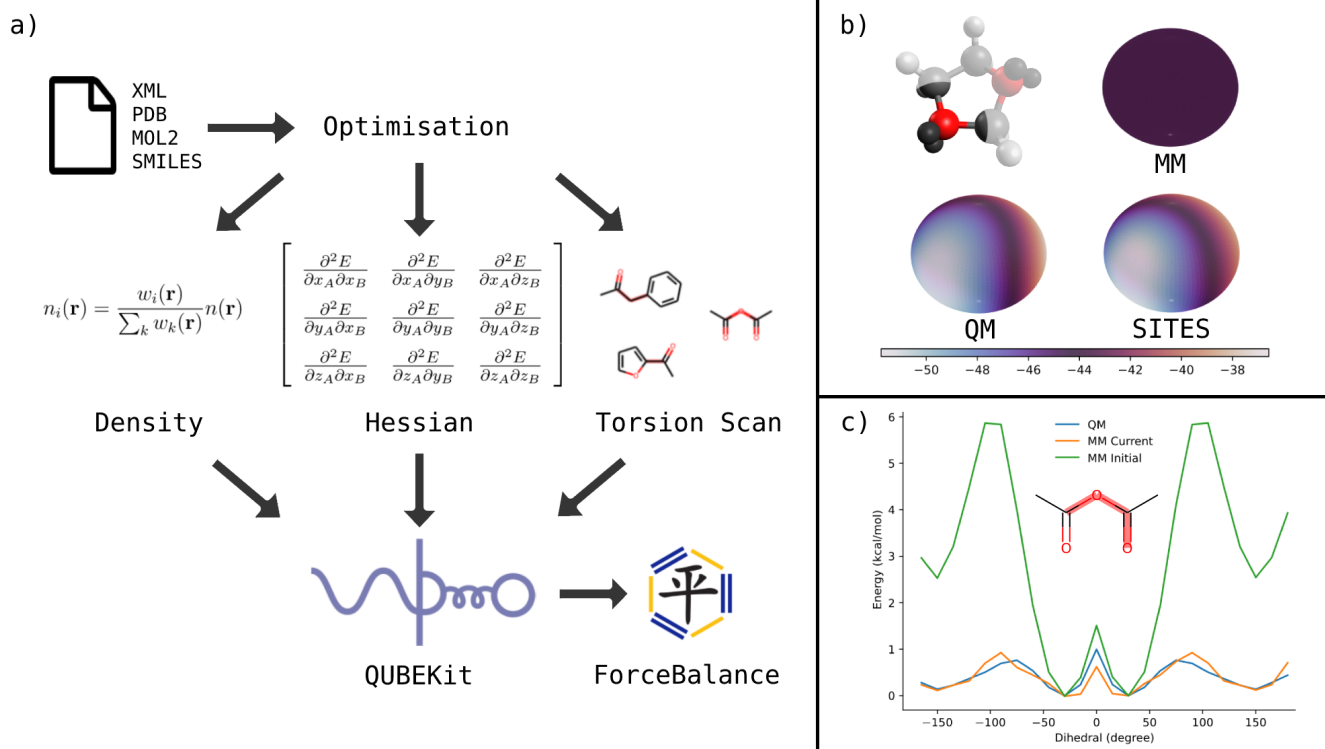


Figure 1: a) QUBEKit workflow. The molecule under study is input and optimised. The total electron density is partitioned into atomic contributions, while Hessian matrices and optional torsion scans are computed. QM measurements are mapped to MM force fields using the QUBEKit software, and (optionally) mapping parameters are optimised using ForceBalance. b) Virtual site derivation for 1,3-dioxolane. (Top right) ESP surface plot (at 1.4x the van der Waals radius) using a single atom-centred point charge (kcal/mol), (bottom left) the QM ESP up to quadrupole order, (bottom right) the MM ESP with the addition of two virtual sites (as shown in top left). c) Example torsion scans before (green) and after (orange) parameter optimisation.

pdb or mol2 file. Initial parameters may be generated from the Open Force Field toolkit,¹¹ Antechamber,⁴⁹ or a supplied XML file. From there, multiple conformers of the molecule coordinates are generated and used in a fast but basic optimisation via TorchANI,⁵⁰ XTB⁵¹ or OpenMM.⁵² Of these conformers, the lowest energy result is taken to the QM optimisation stage. The QM structural optimisation is performed using either PSI4⁵³ or Gaussian,⁴⁸ both run through QCEngine⁵⁴ with an ultra-fine grid. If the first, lowest energy conformer fails to optimise after 50 iterations, the next conformer is passed instead. The QM optimised coordinates are used as input for an electron density calculation, a Hessian matrix calculation and optional 1D torsiondrive calculation(s) (Figure 1(a)). The Hessian matrix is used to calculate all bond and angle parameters (k_r , r_0 , k_θ , θ_0) via the modified Seminario method.³⁷

QUBE uses atoms-in-molecule (AIM) electron density partitioning to obtain all non-bonded parameters of the force field. AIM methods partition the total electron density $n(\mathbf{r})$ into overlapping atomic densities $n_i(\mathbf{r})$ according to:

$$n_i(\mathbf{r}) = \frac{w_i(\mathbf{r})}{\sum_k w_k(\mathbf{r})} n(\mathbf{r}) \quad (2)$$

where the form of the weights, $w_i(\mathbf{r})$, are determined by the choice of AIM method. Common choices include (iterative) Hirshfeld,²⁶ MBIS²³ and density derived electrostatic and chemical (DDEC)^{24,25} partitioning. Following AIM partitioning, atom-centred partial charges are simply assigned by integrating the atomic electron densities:

$$q_i = z_i - \int n_i(\mathbf{r}) d^3\mathbf{r} \quad (3)$$

where z_i is the nuclear charge. It is well-known that partial charges in non-polarisable force fields need to be scaled in some way to account for polarisation in an effective manner in condensed phase simulations. Common approaches include use of bond order corrections or charge scaling factors,^{19,20} but we prefer to use an implicit solvent model with a dielectric constant intermediate between the gas and water phases.^{22,28,55} The final charges thus depend on

the choice of underlying QM method used to compute $n(\mathbf{r})$, the choice of AIM method, and the choice of implicit solvent model and parameters. Previously we computed the electron density using the linear-scaling DFT code, ONETEP,⁵⁶ using the PBE exchange-correlation functional and a non-orthogonal generalised Wannier function basis set.²⁹ A multigrid Poisson solver⁵⁷ was used with a dielectric of 4 to pre-polarise the charges, and an implementation of the DDEC AIM method in ONETEP (similar to DDEC3³⁹) was employed to partition the electron density. As discussed in the Introduction, we now have the infrastructure in place to systematically test the effects of these hyperparameter choices. Table 1 summarises the alternative model protocols that will be investigated here to test the effect of the underlying QM method (models **1a** and **1b**), the implicit solvent parameters (models **2a–2c**) and AIM partitioning scheme (models **3a** and **3b**) on the force field parameters, and physical properties.

In the original QUBE force field, the non-bonded part of the force field includes the Lennard-Jones interaction:

$$V_{LJ}(r_{ij}) = \frac{A_{ij}}{r_{ij}^{12}} - \frac{B_{ij}}{r_{ij}^6} \quad (4)$$

between pairs of atoms i, j . The dispersion coefficient, B_i is derived for the atom in the molecule via the Tkatchenko-Scheffler (TS) relation:²⁷

$$B_i = \left(\frac{V_i^{AIM}}{V_i^{free}} \right)^2 B_i^{free} \quad (5)$$

where V_i^{AIM} is the third radial moment of the partitioned atomic electron density:

$$V_i^{AIM} = \int r^3 n_i(\mathbf{r}) d^3\mathbf{r} \quad (6)$$

The corresponding quantities B_i^{free} and V_i^{free} may be computed and tabulated for free atoms in a vacuum (**Table S1**). To ensure that the Lennard-Jones potential has a minimum at the effective van der Waals radius of the atom-in-molecule, the coefficient of the repulsive term

may be approximated via:²⁸

$$A_i = \frac{1}{2} B_i (2R_i^{AIM})^6 \quad (7)$$

where the AIM radius, R_i^{AIM} , is again found by re-scaling a reference free atom radius:

$$R_i^{AIM} = \left(\frac{V_i^{AIM}}{V_i^{free}} \right)^{1/3} R_i^{free} \quad (8)$$

The R_i^{free} are free parameters to be fit to experiment, which we have found to be crucial if the force field is to reproduce condensed phase properties. Previously, we have fit these via parameter scans to liquid properties (densities and heats of vaporisation).²⁸ However, as discussed in the Introduction, a new set of parameters is required for each set of model hyperparameters, and so we employ here a more automated approach as set out later.

Transforming the A_i and B_i parameters to the σ_i and ϵ_i parameters of eq 1, we obtain (a full derivation is given in **Supporting Information S1.1**):

$$\sigma_i = 2^{5/6} \left(\frac{V_i^{AIM}}{V_i^{free}} \right)^{1/3} R_i^{free} \quad (9)$$

and:

$$\epsilon_i = \frac{B_i^{free}}{2 \left(2R_i^{free} \right)^6} \quad (10)$$

Throughout this study, the Lorentz-Berthelot combination rules are used to derive ϵ_{ij} and σ_{ij} from the corresponding atomic quantities. Eq 9 results in atoms with more diffuse electron density having a larger Lennard-Jones σ_i parameter but, as correctly pointed out previously,³⁵ all atoms of the same element have the same Lennard-Jones well depth (eq 10). In that same study, the authors proposed that each atom should have a unique effective ionisation energy, derived from the partitioned atomic electron density, which affects the scaling relationship in eq 5.³⁵ An additional consideration is that the dispersion coefficient (B_i) should be viewed as an effective interaction, taking into account not only the dipole-

dipole, but also higher-order (dipole-quadrupole etc) contributions to dispersion. It has been shown that physics-based derivations of B_i tend to therefore be lower than the effective coefficients in common MM force fields.^{30,40} To take both of these factors into account, we test here the effect of introducing additional flexibility in the definition of the dispersion coefficient:

$$B_i = \alpha \left(\frac{V_i^{AIM}}{V_i^{free}} \right)^{2+\beta} B_i^{free} \quad (11)$$

where α and β are global fitting parameters. This change has no effect on σ_i , but ϵ_i is now dependent on the diffuseness of the atomic electron density, measured by V_i^{AIM} (a full derivation is given in the **Supporting Information S1.2**):

$$\epsilon_i = \frac{\alpha \left(\frac{V_i^{AIM}}{V_i^{free}} \right)^{\beta} B_i^{free}}{2 \left(2R_i^{free} \right)^6} \quad (12)$$

Note that β can be positive or negative. Thus, as summarised in Table 1, we can set $\alpha = 1$ and $\beta = 0$ to retain the original QUBE force field model (e.g. model **0**), or optimise α and β alongside R_i^{free} (e.g. model **4b**). The effects of these choices are discussed later.

An additional choice in force field design is whether to include Lennard-Jones parameters on polar hydrogen atoms, or to effectively redistribute them onto the neighbouring heavy atom²⁸ (see **Supporting Information S1.3**). Different choices are made for example in the design of the TIP3P and CHARMM modified water model,⁵⁸ while there is precedent for improved agreement with liquid data using Lennard-Jones parameters on polar hydrogen atoms.⁵⁹ While separate mapping parameters have been previously proposed for polar and non-polar hydrogen atoms, the effect on accuracy has not been directly explored,³⁵ and so we add this comparison to our list of model protocols to be tested (Table 1, model **4a**).

A final consideration to be made when deriving the non-bonded parameter set is how to treat atoms with significant anisotropy in their electron density, such that an atom-centred point charge gives a poor approximation to the full QM electrostatic potential at the surface

of the molecule. Such situations are commonly treated in MM force fields using an off-centre charge (or “virtual site”), as used in, for example, the TIP4P water model⁶⁰ and various treatments of lone-pairs and σ -holes in common force fields.^{61–63} In keeping with our general force field philosophy, we previously proposed a method to derive the positions and magnitudes of off-site charges, where required, directly from the partitioned QM electron density. Full details are given elsewhere,²⁹ but in brief the charges are positioned to minimise the quantity:

$$\Phi = \frac{1}{n} \sum_{i=1}^n |V_{QM}^i - V_{MM}^i| \quad (13)$$

where V_{QM}^i is the electrostatic potential generated by the *partitioned* atomic electron density at a sample point i , V_{MM}^i is the corresponding quantity calculated using the point charge model, including any off-centre charge(s), and n is the total number of sample points (located between 1.4 to 2.0 times the van der Waals radius of the atom). By using the atomic, rather than the molecular, electron density, the method scales easily to larger molecules. To limit the search space, off-site charges were restricted to positions determined by the symmetry of the atom’s bonding environment.²⁹ For example, for halogens bonded to a single atom, the search direction is along the bond vector, and for oxygen bonded to two neighbours, the search direction is along the bisector of the two bonds. Similar arguments can be made for positioning two virtual sites.²⁹

A similar protocol is followed in the current work. Key differences are discussed here and full implementation details are provided in the **Supporting Information Section S2**. It is desirable to perform the optimisation of eq 13 directly in QUBEKit, but storage and imports of the full atomic electron density is inefficient. Instead, V_{QM}^i is now reconstructed from the QM atomic multipole moments, up to quadrupole order, which are readily extracted from either Chargemol (using the DDEC AIM method)^{24,25} or the PSI4 software (using the MBIS AIM method).⁵³ For atoms with isotropic electron density, the atomic multipoles are low, and the QM electrostatic potential is well-described by an atom-centred point charge (monopole). For atoms with anisotropic electron density (defined here as $\Phi > 1$ kcal/mol),

eq 13 is minimised with respect to the virtual site charge and position, using the Scipy python package.

Figure 1(b) shows example surface plots of the electrostatic potential energy (ESP) at 1.4x the van der Waals surface of an oxygen atom in 1,3-dioxolane. The bottom-left plot shows the QM-calculated ESP, up to quadrupole order. The top-right plot shows the ESP using an atom-centred monopole, which simply gives a uniform, averaged ESP over the entire surface of the atom. The difference between these two plots at every point allows the calculation of the error function, eq 13. The bottom-right plot shows an approximation of the QM ESP, constructed with the addition of just two virtual sites, which are able to recover the expected anisotropy in the ESP. Table 1 shows the experiments performed in what follows to test the effect of including virtual sites on our force field model accuracy (models **5a–5e**).

In previous works, new non-bonded force field parameters tend to be investigated in conjunction with the bonded parameters of standard transferable force fields, which may lead to incompatibility given the close inter-dependency between parameter types. Here, the benefit of assembling the force fields using the QUBEKit package, is that torsion parameters for rotatable bonds may be re-fit in a simple, automated manner for compatibility with the rest of the bonded and non-bonded parameters. Using a new interface between QUBEKit, ForceBalance³⁶ and TorsionDrive,⁶⁴ dihedral parameters of any freely rotatable bonds were optimised via a least squares minimisation of the root mean square error (RMSE) between QM and MM torsional potential energy surfaces. We use the same procedure for parameter fitting as the recent Open Force Field Parsley force field,¹¹ and full details are provided in **Supporting Information Section S3.1**. A new feature in this work is that we have modified the ForceBalance code to allow parameter fitting in the presence of local coordinate virtual sites. These virtual sites then inherit their exceptions and exclusions from the parent atom meaning they interact with the same rules as the parent. The code may be obtained from conda-forge or github (<https://github.com/leeping/forcebalance>), and is available from

version 1.9.0 onward. Figure 1(c) demonstrates an example of bespoke dihedral parameter fitting for the molecule acetic anhydride, which has two virtual sites on the central oxygen atom. With the initial parameter set (which is incompatible with the new Lennard-Jones parameters and virtual sites), the RMS error is more than 3 kcal/mol, but following fitting it falls to just 0.12 kcal/mol. In fact, the average RMSE across all molecules in the 51 molecule test set (a total of 117 scans) after fitting was just 0.13 kcal/mol.

ForceBalance Workflow

As outlined in Figure 1(a), for each choice of force field design protocol (Table 1), the ForceBalance software³⁶ was employed to optimise the set of fitting parameters. These parameters were the R_i^{free} parameters of H, C, N and O (eq 8), and optionally the dispersion rescaling parameters α and β (eq 11). Where Lennard-Jones parameters were included on polar hydrogen atoms, it was found that a separate R_i^{free} parameter was required for these atoms. Optimisation was performed against the experimental liquid densities and enthalpies of vapourisation of a training set of 15 molecules. The training set comprised molecules containing H, C, N and O only, and a range of functional groups (**Supporting Information Section S4**). It was identical to that used in the training of the RESP2 charge fitting procedure,²² except that one molecule (N1(C)CCOCC1) was replaced by CN(C)C=O to ensure that sampling complex potential energy surfaces did not influence the fit. ForceBalance uses the reference data to produce an objective function, which is minimised with a non-linear optimisation algorithm through numerical differentiation of physical properties with respect to the fitting parameters. Further details are given in the **Supporting Information Section S3.2**. To create the inputs required by ForceBalance, QUBESKit automatically combines the force fields of the individual molecules, along with algebraic expressions that transform the fitting parameters into the σ_i and ϵ_i parameters of the force field (see, for example, eqs 9 and 10), into a single xml file. The full training set, experimental data, example ForceBalance input files, and tutorial are provided in the **Supporting Information**.

Following the completion of ForceBalance training runs, the final fitting parameters are used to re-calculate the Lennard-Jones and re-fit the torsion parameters in QUBEKit. Training set accuracy was then confirmed by re-computing the liquid properties of the 15 molecules using QUBEBench, which is an interface with OpenMM for the benchmarking of QUBE force fields (**Supporting Information Section S3.3**). QUBEBench and ForceBalance data deviated by just 0.002 g/cm³ and 0.11 kcal/mol, on average, for density and heat of vaporisation calculations, respectively, indicating good convergence of the fit. Finally, four of the most encouraging force field protocols were taken forward to testing. Here, QUBEKit was used to fit force fields for a test set comprising a further 51 molecules taken from a previous study (**Supporting Information Section S4**).²² Physical properties were calculated using QUBEBench, and compared with experiment, to confirm transferability of the designed protocols.

Results

Training a range of force field models identifies patterns in accuracy.

Using our new interface between the QUBEKit engine for QM-to-MM force field parameter mapping and the ForceBalance software for parameter tuning, we have trained a total of 15 different force field protocols to investigate the effects of choices made in force field design on model accuracy. Descriptions of the force field protocols are listed in full in Table 1, grouped roughly by the type of hyperparameter being investigated.

The first groupings investigate the choice of underlying QM methods used to compute the total electron density for AIM partitioning and subsequent non-bonded parameter derivation (the same QM method is also used for Hessian matrix calculation and dihedral scans to obtain the bonded parameters). As mentioned, our first implementation of QUBEKit²⁹ interfaced with the ONETEP DFT code,⁵⁶ which uses the PBE exchange-correlation functional, for non-bonded parameter derivation. However, this new version of QUBEKit is interfaced, via

Table 1: Summary of changes across force field models. Model **0** is designated as the default protocol, and changes from the default in subsequent models are highlighted in bold. Training set mean unsigned error (MUE) in density (g/cm^3) and heat of vaporisation (kcal/mol), relative to experiment, are displayed for each force field protocol.

| Model | QM Method | Solvent | V-Sites | AIM | LJ | MUE (ρ , ΔH_{vap}) |
|-----------|---|---------------------|------------|--------------|-----------------|-----------------------------------|
| 0 | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | No | DDEC6 | TS | 0.0274, 0.766 |
| 1a | B3LYP-D3(BJ)/DZVP | IPCM, $\epsilon=4$ | No | DDEC6 | TS | 0.0271, 0.726 |
| 1b | HF/6-31G(d) | None | No | DDEC6 | TS | 0.0248, 0.731 |
| 2a | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=2$ | No | DDEC6 | TS | 0.0296, 0.772 |
| 2b | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=10$ | No | DDEC6 | TS | 0.0237, 0.783 |
| 2c | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=20$ | No | DDEC6 | TS | 0.0285, 0.678 |
| 3a | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | No | DDEC3 | TS | 0.0213, 0.475 |
| 3b | B3LYP-D3(BJ)/DZVP | Chloroform | No | MBIS | TS | 0.0159, 0.578 |
| 4a | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | No | DDEC6 | H0 | 0.0235, 0.896 |
| 4b | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | No | DDEC6 | α, β | 0.0206, 0.587 |
| 5a | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | Yes | DDEC6 | TS | 0.0190, 0.441 |
| 5b | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | Yes | DDEC6 | α, β | 0.0209, 0.244 |
| 5c | $\omega\text{B97X-D}/6\text{-}311++\text{G(d,p)}$ | IPCM, $\epsilon=4$ | Yes | DDEC3 | TS | 0.0163, 0.357 |
| 5d | B3LYP-D3(BJ)/DZVP | IPCM, $\epsilon=4$ | Yes | DDEC6 | TS | 0.0141, 0.450 |
| 5e | B3LYP-D3(BJ)/DZVP | Chloroform | Yes | MBIS | α, β | 0.0175, 0.480 |

Table 2: Values of the fitting parameters following ForceBalance optimisation for each of the force field protocols. The R_i^{free} for each element are in Å, and α and β are dimensionless.

| Model | C | N | O | H | polar H | α | β |
|-----------|-------|-------|-------|-------|---------|----------|---------|
| 0 | 2.008 | 1.765 | 1.499 | 1.738 | 1.083 | — | — |
| 1a | 1.999 | 1.740 | 1.489 | 1.752 | 1.111 | — | — |
| 1b | 2.042 | 1.676 | 1.501 | 1.737 | 1.218 | — | — |
| 2a | 2.004 | 1.708 | 1.464 | 1.744 | 1.107 | — | — |
| 2b | 2.026 | 1.751 | 1.521 | 1.724 | 1.119 | — | — |
| 2c | 2.025 | 1.756 | 1.503 | 1.733 | 1.132 | — | — |
| 3a | 2.051 | 1.740 | 1.590 | 1.670 | 1.126 | — | — |
| 3b | 2.068 | 1.681 | 1.599 | 1.753 | 1.404 | — | — |
| 4a | 2.021 | 1.604 | 1.550 | 1.719 | — | — | — |
| 4b | 2.074 | 1.742 | 1.481 | 1.760 | 1.154 | 1.301 | 0.465 |
| 5a | 1.994 | 1.706 | 1.558 | 1.738 | 1.279 | — | — |
| 5b | 2.035 | 1.722 | 1.574 | 1.731 | 1.294 | 1.221 | 0.489 |
| 5c | 2.042 | 1.740 | 1.630 | 1.687 | 1.274 | — | — |
| 5d | 2.013 | 1.680 | 1.558 | 1.732 | 1.442 | — | — |
| 5e | 2.043 | 1.693 | 1.680 | 1.680 | 1.464 | 0.999 | 0.491 |

QCEngine,⁵⁴ with both Gaussian09⁴⁸ and PSI4,⁵³ which gives us access to a wide range of alternative quantum chemistry methods. If using PSI4, the electron density partitioning is performed using MBIS.²³ If using Gaussian09, Chargemol is used to calculate the atom centred charges via either the DDEC3 or DDEC6 AIM schemes. In all cases, the AIM method returns atom-centred point charges, the third radial moment of the atomic electron density (for Lennard-Jones parameter derivation), and multipole moments of the atomic electron density (for virtual site derivation).

QM Method

First, we investigate the use of two relatively high level QM methods. The ω B97X-D/6-311++G(d,p) method is used for bonded parameter derivation in earlier QUBE force fields,^{29,65} as well as the OPLS-AA/M protein force field.⁶⁶ The B3LYP-D3(BJ)/DZVP method is used for bonded parameter fitting in the Parsley force field¹¹ and has been shown to give a good compromise between accuracy and computational expense for gas phase conformational energetics of small organic molecules.⁶⁷ It has been shown that the choice of DFT functional is not too critical in the reproduction of molecular electrostatic potentials, while larger basis sets afford some improvements in accuracy.²² However, relatively little is known about how this choice translates to force field accuracy. Runs **0** and **1a** reveal that this choice is not critical, with nearly identical training set accuracy obtained for both liquid densities (0.027 g/cm³) and heats of vaporisation (0.77 vs 0.73 kcal/mol). It should be emphasised that this does not mean that the non-bonded parameters themselves are identical. Table 2 shows the final R_i^{free} parameters after ForceBalance training. There is a notable reduction in the radius of the N atom, using the B3LYP-D3(BJ)/DZVP method, and corresponding increase in both the polar and non-polar H atoms. This backs up the recent assertion that improving the accuracy of atomic charges is unlikely to improve calculation accuracy without a corresponding optimisation of Lennard-Jones parameters.²²

For comparison, in run **1b**, we employed the HF/6-31G(d) method, which has been used

historically for the fitting of force field ESP charges.²¹ The argument goes that this method produces a fortuitous over-polarisation of the electron density in the gas phase to yield charges that are suitable for condensed phase modelling, and so we perform these electron density calculations in vacuum (Table 1). The results are very similar to the two higher level QM methods in implicit solvent, but perhaps this is not too surprising, given the success of many force fields that are built from the HF/6-31G(d) method. However, as we will show, we can now do better with improved physics-based protocols.

Implicit Solvent Model

In previous versions of QUBEKit, we have computed the electron density using a minimal parameter solvent model, with the solute cavity defined by an isosurface of the electron density.⁵⁷ We argued that a solvent dielectric $\epsilon = 4$ is appropriate, as it leads to charges that are polarised midway between vacuum and a dielectric medium of $\epsilon = 78$ (i.e. water).²⁸ Gaussian09 provides the IPCM implicit solvent model,⁶⁸ which also builds the solute cavity from an isosurface of the electron density, and so we use this model as our default here. Models **2a–2c** compare the effect of tuning the solvent dielectric in the range $\epsilon = 2$ –20. Again, very little change in force field accuracy is observed, with density errors in the range 0.02–0.03 g/cm³ and heat of vaporisation errors in the range 0.6–0.8 kcal/mol. This is broadly in agreement with previous observations,²² which showed a strong dependence of condensed phase properties on the polarity of the charge model for a fixed set of Lennard-Jones parameters, but much less sensitivity when the Lennard-Jones parameters are tuned for consistency with the charges. Table 2 shows that the R_i^{free} parameters in run **2a** tend to be lower than the others, which indicates that the minimum of the Lennard-Jones interaction should be closer to the atom for the less-polarised charge model. Compared to vacuum, we found that molecular dipole moments are scaled by a factor of 1.10x ($\epsilon = 2$) and 1.22x ($\epsilon = 20$), on average, indicating that a suitable range of condensed phase polarisation has been captured in these force field models by tuning the dielectric of the background implicit

solvent (**Figure S4**).

AIM Partitioning Method

Our previous version of QUBEKit relied on an implementation of the DDEC AIM approach in the ONETEP DFT code.³⁹ The DDEC AIM method is described elsewhere,²⁴ but in short it aims to construct the optimal weighting factors of eq 2 that result in approximately spherical atomic electron densities, so that its multipole expansion converges rapidly, whilst ensuring that the resulting charges are chemically reasonable. Alongside the observation that DDEC charges show excellent transferability between different conformations of the same molecule,²⁵ these properties make the DDEC AIM approach very promising for flexible force field design.

In the previous sections, we employed the latest DDEC6 AIM approach for the first time in organic molecule force field design, through the interface between Gaussian09 and Chargemol. In model **3a**, we investigate the effect of switching to the older DDEC3 approach. The accuracy is actually slightly better for DDEC3, with a decrease in both density and heat of vaporisation errors, compared to the baseline model **0**. DDEC6 has several methodological improvements, which are summarised elsewhere²⁴ and have been shown to result in more robust convergence, lower computation times, lower ESP errors, and higher transferability, when compared to DDEC3.^{24,25} However, most of these advantages will not be apparent in the small, relatively rigid, organic molecules, which lack buried atoms and are simulated here under standard conditions.

MBIS is similar in idea to the Hirshfeld family of AIM methods.²³ In this method the total electron density is partitioned onto a minimal set of s-type Slater functions, whose parameters are fit to the input molecular electron density by minimising the Kullback-Leibler divergence. Like DDEC, it has been shown that the MBIS-derived charges reproduce the QM ESP to good accuracy and are robust to small conformational changes, while MBIS can also be used to derive Tkatchenko-Scheffler based dispersion coefficients.²³ MBIS charges, along

with all AIM properties (atomic multipoles up to quadrupole order and radial moments) that we require for force field derivation have recently been implemented in the PSI4 quantum chemistry package,⁵³ which we can now access through our latest interface with QCEngine.⁵⁴ Model **3b** thus investigates the use of MBIS as the AIM partitioning method in force field derivation. For technical reasons, it was necessary to use the B3LYP exchange-correlation functional, and the IEFPCM implicit solvent model (with a chloroform solvent to mimic a dielectric of approximately 4),⁶⁹ for the underlying QM calculations. Again this combination of methods gave very good accuracy on our training set, with errors in the density and heat of vaporisation of 0.016 g/cm³ and 0.58 kcal/mol, respectively.

Given the expected benefits of DDEC6 in larger, more flexible molecules, we retain this method as our standard approach, but do investigate both DDEC3 and MBIS further in what follows. It is gratifying that a range of AIM methods have been shown to be promising for flexible force field design, and are available through our QUBEKit interface with QCEngine.⁵⁴

Lennard-Jones Parameters




Figure 2: Correlation between Parsley and QUBE (model **5b**) summed molecular dispersion coefficients for each molecule in the test set ($B^{mol} = \sum_{i \in mol} B_i$).

As discussed in the Methods section, the default QUBE protocol for deriving the attractive part of the Lennard-Jones potential (B_{ij} in eq 4) from QM tends to lead to dispersion coefficients that are lower than those used in common effective force fields⁴⁰ and energy

well depths that are constant across atoms of the same element.³⁵ Model **4b** in Table 1 investigates the effects of relaxing this protocol, by introducing two new global, variable parameters (α and β) into the fit, thus moving away from the default Tkatchenko-Scheffler (TS) approach. As expected, a reduction in training set errors is achieved, relative to the default model **0**, with these extra fitting parameters. Table 2 shows that this is achieved with $\alpha > 1$ and $\beta > 0$, both factors that act to increase the well depth for atoms with more diffuse electron density (eq 12). Figure 2 plots the dispersion coefficients (summed over all atoms in the molecule) for all molecules in the larger test set, using the identical Lennard-Jones scheme (from model **5b**, see later), against the corresponding quantities extracted from the Parsley force field. Unlike previous protocols (see our previous work,⁴⁰ and **Section S5.2**), there is a very good correlation between the two force fields, albeit with a consistent offset. While the Parsley parameters are extracted from a library, which in turn has been fit to condensed phase properties, and ours are derived from QM calculations specifically for the molecule under question, the convergence of the two approaches is encouraging.

In addition, we investigate the question of whether to include Lennard-Jones parameters on polar hydrogen atoms, or to set them to zero and effectively absorb them onto the neighbouring heavy atom (**Section S1.3**).²⁸ In Table 1, we show that force field protocol **4a**, which has no Lennard-Jones parameters on polar hydrogens (denoted ‘H0’), has a similar density error, but higher heat of vaporisation error (0.90 kcal/mol) compared to the other protocols. Interestingly, this model has a much lower R_i^{free} for nitrogen than the other models, but the low density error indicates that this does not propagate through to too short hydrogen bonding distances in condensed phase simulations. Nevertheless, we conclude that explicitly including Lennard-Jones parameters on polar hydrogen atoms has accuracy benefits, and we continue to use this approach in the following force field models.

Virtual Sites

Figure 1(c) shows an example benefit of including just a small number of virtual sites on a molecule in terms of the reproduction of the QM electrostatic potential (in this case up to quadrupole order). In fact, considering a larger test set of 51 molecules, we found that 33 molecules had electron density anisotropy above our set threshold (1 kcal/mol, eq 13) for at least one atom. As expected from our previous work,²⁹ these virtual sites were found almost exclusively on oxygen and nitrogen atoms. Using atom-centred point charges only, the average ESP error on these atoms was 1.87 kcal/mol, which fell to 0.53 kcal/mol after virtual site fitting. Further examples are given in the next section and a full list of molecules and ESP errors are given in the **Supporting Information**.

However, a more pertinent question is whether a force field model that more accurately captures the QM electrostatic potential also results in more accurate prediction of condensed phase properties. We have therefore trained a series of force field protocols that include virtual sites in the description of the electrostatics (models **5a–5e**). Table 1 summarises the accuracy on the training set. As opposed to some of the other design choices, we now see significant accuracy gains for these models, with predictions in liquid densities in the range 0.014–0.021 g/cm³ and heats of vaporisation between 0.24–0.48 kcal/mol. The lowest density error is observed when virtual sites are used in combination with the B3LYP-D3(BJ)/DZVP QM method (model **5d**), and the lowest energy errors when used with rescaled Lennard-Jones interactions (model **5b**). Although there is no significant trend in the fit R_i^{free} parameters on the heavy atoms for these force field protocols (Table 2), interestingly there is a consistent shift of the van der Waals radii of the polar hydrogen atoms to higher values (1.27–1.46 Å) when used in combination with virtual sites. This makes intuitive sense since these atoms are likely to be involved in hydrogen bonds with anisotropic polar atoms, like oxygen and nitrogen. The virtual sites will move the centre of electron density closer to the van der Waals surface of the heavy atom, meaning that the Lennard-Jones repulsion should increase to compensate the increase in electrostatic attraction. The improved condensed phase prop-

erties seem to indicate that this is a more physically-reasonable balance than the traditional atom-centred point charge approach.

Relative training set accuracy is retained in the test set

Table 3: Test set accuracy. MUE in density (ρ) and heat of vaporisation (ΔH_{vap}) for 51 molecules in the test set. 95% confidence intervals are included.

| Force Field | ρ (g/cm ³) | ΔH_{vap} (kcal/mol) |
|-----------------|--|---|
| Parsley | 0.038 ^{0.048} _{0.028} | 1.18 ^{1.53} _{0.87} |
| Model 0 | 0.033 ^{0.043} _{0.024} | 1.53 ^{1.87} _{1.17} |
| Model 1a | 0.041 ^{0.052} _{0.030} | 1.33 ^{1.58} _{1.11} |
| Model 5b | 0.036 ^{0.046} _{0.027} | 0.69 ^{0.86} _{0.51} |
| Model 5d | 0.031 ^{0.041} _{0.022} | 1.02 ^{1.33} _{0.73} |

To test the accuracy of the designed force fields on molecules outside the training set, we ran simulations of the condensed phase properties of a separate test set of 51 molecules containing H, C, N and O atoms only. The force fields brought through for study were models **0**, **1a**, **5b** and **5d**. The first two were chosen as our simplest (but also amongst the least accurate) protocols, differing only by the choice of underlying QM method. In both cases, we see a drop in overall accuracy, compared to the training set, which is to be expected for the larger, more complex molecules found in the test set (**Figure S3**). Again, we do not observe a strong dependence of the results on the choice of QM method, and so the computationally less expensive model **1a**, based on the B3LYP–D3(BJ)/DZVP method, should make a good compromise force field. For comparison, we have also run the test set using the Open Force Field Parsley force field¹¹ using identical protocols. The accuracy in the simulated densities (0.038 g/cm³) is similar to our first two protocols, and the error in the heat of vaporisation (1.18 kcal/mol) is slightly lower. The Parsley force field is based around the AM1-BCC charge scheme, with Lennard-Jones parameters largely retained from the SMIRNOFF99Frosst force field.⁷⁰ Thus, there is scope for further refinement of the Lennard-Jones parameters within the current functional form, albeit with many more adjustable parameters than we use here (the five R_i^{free} parameters in Table 2).

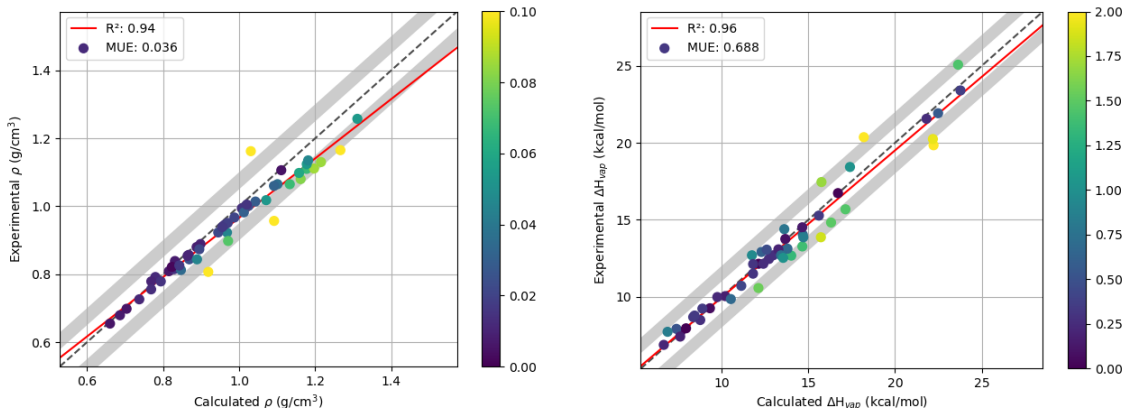


Figure 3: Correlation between experimental and calculated (model **5b**) physical properties of liquids in the test set.

Two additional force field protocols were chosen for further investigation based on low training set errors. Model **5d** uses the B3LYP–D3(BJ)/DZVP method, and virtual sites to model anisotropic electron density. In this case, we see moderate improvement in test set accuracy for both density and heat of vaporisation. Model **5b** uses our default quantum chemistry method, virtual sites and re-scaled Lennard-Jones interactions. With approximately the same density error, the error in the heat of vaporisation is significantly reduced to around 0.7 kcal/mol (Figure 3). Only two compounds have heat of vaporisation errors exceeding 2 kcal/mol, and these are two relatively long chain molecules containing hydroxyl groups (SMILES: OCCCCCO and CC(C)(O)CCC(C)(C)O).

It is difficult to pinpoint reasons for improvement in accuracy, since the physical properties are derived from many competing effects. However, Figure 4 shows some of the force field parameters for molecules that show differences in performance between our best model (**5b**), and Parsley. Figure 4(a) shows 1,3-benzodioxolane, which has $\Delta H_{vap}^{exper} = 13.1$, $\Delta H_{vap}^{QUBE} = 13.8$ and $\Delta H_{vap}^{parsley} = 15.3$ kcal/mol. The net charge and Lennard-Jones ϵ parameter is similar for both QUBE and Parsley, but the former has the charge roughly evenly spread over the atom centre and two virtual sites, showing the benefits of modelling anisotropy in the electron density. Figure 4(b) shows heptane, and in this case both force fields are in excellent agreement with each other and experiment ($\Delta H_{vap}^{exper} = 8.8$, $\Delta H_{vap}^{QUBE} = 8.4$ and




Figure 4: A comparison of selected force field parameters between QUBE (model **5b**), and Parsley. Charges on the atom centre and virtual sites (if applicable) are shown, along with the Lennard-Jones ϵ parameter (kcal/mol).

$\Delta H_{vap}^{parsley} = 8.7$ kcal/mol). There are some small differences in charges on the terminal methyl hydrogen atoms, which add up to a large difference in charge on the carbon atoms, but this does not seem to significantly affect the liquid properties. Figure 4(c) shows methyl isocyanate, for which QUBE model **5b** is in much better agreement with experiment than Parsley ($\Delta H_{vap}^{exper} = 6.9$, $\Delta H_{vap}^{QUBE} = 6.7$ and $\Delta H_{vap}^{parsley} = 11.0$ kcal/mol). This seems to be a general issue with the atom-centred fixed charge force fields, since using our model **0**, we obtain $\Delta H_{vap}^{QUBE} = 10.1$ kcal/mol. Indeed, QUBE model **5b** has a smaller net charge on the N atom of methyl isocyanate than Parsley, but some of the charge is spread onto a virtual site. Finally, Figure 4(d) shows piperidine, for which slightly improved agreement with experiment is again obtained ($\Delta H_{vap}^{exper} = 9.2$, $\Delta H_{vap}^{QUBE} = 8.9$ and $\Delta H_{vap}^{parsley} = 10.5$ kcal/mol) with a smaller net charge and a virtual site on the N atom.

We can compare the accuracy of the protocols developed here with other force fields tested on the same or similar test sets. Using the RESP2 charge assignment method, which is based on fitting to the QM electrostatic potential in implicit solvent, with optimised Lennard-

Jones parameters, the density and heat of vaporisation errors on an identical test set are 0.024 g/cm³ and 1.67 kcal/mol.²² Thus, substantial improvement in energetics is obtained in the current study, at the expense of a slight degradation in liquid density prediction. In that study, only a subset of ten Lennard-Jones parameters (the well depth and radius for each of C, O, N, H and polar H) were optimised for efficiency, though this is still more parameters than our most complex protocol (model **5b** has seven adjustable parameters). Kantonen et al have used a similar QM-to-MM parameter mapping approach. They derive Lennard-Jones parameters from the MBIS partitioned atomic electron densities, with two fitting parameters (used to map the electron density decay constants onto the Lennard-Jones σ and ϵ) for each of C, O, N, H, and polar H. These QM-derived Lennard-Jones parameters are used alongside the AM1-BCC charge model, and GAFF bonded parameters. On a test set of 23 small organic molecules, they report MUEs of 0.027 g/cm³ and 1.1 kcal/mol in density and heat of vaporisation, respectively, which is similar in performance to our model **5d** protocol. Visscher and Geerke have used an AIM-based QM-to-MM mapping scheme, using both C_6 and C_8 dispersion coefficients, in combination with an ESP-based charge scheme and a charge-on-spring polarisable model.³⁴ Our strategy of fitting a small number of R_i^{free} parameters was also employed there, but no off-site charges were used. Encouraging root-mean-square deviations from experiment of 0.024 g/cm³ and 0.39 kcal/mol in ρ and ΔH_{vap} were reported on a combined train/test set of 49 small organic molecules.³⁴ With our automated methods for force field derivation and training, it appears that further accuracy gains can be expected by moving in future to polarisable and beyond-Lennard-Jones models.

Discussion and Conclusions

Force field design is a lengthy process, often requiring large teams of researchers working over a period of many years. Design decisions are typically made early in the process, and any

inaccuracies stemming from these assumptions propagate through to the final force field. Even the minimal Parsley force field has a set of 35 Lennard-Jones parameters,¹¹ and so large training sets are required to cover sufficient chemical space. Repeatedly performing the full training process to investigate alternative design decisions would substantially slow down the time to science. Furthermore, as we inevitably move to more complex and accurate functional forms, the size of the parameter space and danger of over-fitting will only become worse.

To solve this problem, we and others have advocated making use of QM-to-MM mapping procedures to significantly reduce the number of parameters that require fitting to experimental physical properties, while retaining a small amount of empiricism to ensure accurate condensed phase properties. Though there are isolated examples in the literature, QM-to-MM mapping potentials are not widely available for general use in deriving full MM force fields for organic molecules. In QUBEKit, and its interface with ForceBalance, we provide here a means to train, test and apply these force fields, deriving parameters for all (including bonded) terms in the force field from QM. By using AIM-based charges, Lennard-Jones parameters and virtual sites, we have a consistent set of non-bonded parameters derived from a single set of atomic electron densities, available through widely-used quantum chemistry software packages. By coupling these parameter sets through QUBEKit, we can derive bonded parameters that are specific to the molecule under study, and fully consistent with the non-bonded parameter set. We estimate that for the 51 molecules in the test set, a total of around 20 k parameters are derived from QM, while only seven parameters are fit to experiment (in model **5b**). Despite this minimal fitting approach, the model achieves highly competitive accuracy on a challenging test set of liquid property data.

We have used this design and train cycle to answer a number of fundamental questions about force field design protocols. It was shown that the details of the underlying QM method and implicit solvent model have negligible effect on the accuracy of condensed phase modelling with the designed force fields. That is not to say that any set of charges gives

identical results, but rather that as long as the charges and Lennard-Jones parameters are co-optimised, the overall accuracy of the force field is insensitive to the particular choice of model. In this case, future efforts should focus on the cheapest QM methods that continue to give reliable results. It is well known that there is no unique method to partition the total molecular electron density into atom-centred basins. We have therefore investigated two variants of the DDEC AIM method, and the MBIS approach. Surprisingly, the older DDEC3 method showed improvements over DDEC6, but given the improvements in the latter in robustness and transferability of the charges, we continue to use this approach. The MBIS method provides a useful alternative AIM approach, available through the open source PSI4 software package. Once atomic electron densities are assigned, the details of the Lennard-Jones parameterisation also makes a difference to force field accuracy. Based on our data, we advocate setting non-zero Lennard-Jones parameters on polar hydrogen atoms, and re-scaling the strength of the QM-derived dispersion parameter. The latter approach gives not only an improvement in accuracy, but brings QM-derived and Parsley empirical Lennard-Jones parameters into closer agreement. Finally, consistent improvement in both training and test set accuracy is achieved by introducing a small number of off-site charges to model anisotropy in the QM electron density. It is encouraging that improving the underlying physics of the force field model, with minimal costs to run time, results in such improvements.

In the current study, we have focused on liquid density and heat of vaporisation as measures of force field accuracy. In future, additional properties, such as enthalpies of mixing,⁷¹ could be used to extend the confidence in the range of application of the force fields. More complex properties, such as free energies of hydration or even protein-ligand binding could be included (at least for testing), but would require infrastructure improvements to allow for virtual sites. The current study gives confidence that this investment of time would be worthwhile.

A notable feature of the current study is the number of force field models that can be

rapidly designed. We have generated here 15 different models, with varying assumptions, which can be contrasted with the handful of general purpose small molecule force fields that are otherwise available. This opens up the possibility of new research into consensus property predictions using force field ensembles, which has been shown to be advantageous, for example, in recent protein-ligand binding free energy studies.⁷²

More broadly, we envisage QM-to-MM mapping potentials, such as these, providing synergy with traditional force fields. For example, our conclusions concerning force field design protocols can be fed into large-scale fitting efforts, such as the Open Force Field Initiative, to focus efforts in areas where accuracy improvements are expected. This will be particularly important as the community moves towards more advanced functional forms, for which the number of parameters to be fit will only increase. With the advent of machine learning models for parameterising force fields,⁷³ the force fields developed here could be used to provide regularisation of the parameter fits, to avoid unphysical predictions by the models. Or conversely, machine learning models could be trained to reproduce the outputs of our QM-to-MM mapping procedures to substantially reduce the computational cost of the parameterisation stage.⁷⁴ All software required to derive the designed force fields is freely available at <https://github.com/qubekit/QUBEKit>.

Acknowledgement

C.R. is grateful for International Partnership Seed Funding from Newcastle University. D.J.C. and J.H. acknowledge support from a UKRI Future Leaders Fellowship (grant MR/T019654/1). We thank the Open Force Field Initiative and PSI4 developers for scientific software support. This work made use of the facilities of the N8 Centre of Excellence in Computationally Intensive Research (N8 CIR) provided and funded by the N8 research partnership and EPSRC (grant EP/T022167/1).

Supporting Information Available

References

- (1) Dawson, J. A.; Canepa, P.; Famprakis, T.; Masquelier, C.; Islam, M. S. Atomic-Scale Influence of Grain Boundaries on Li-Ion Conduction in Solid Electrolytes for All-Solid-State Batteries. *Journal of the American Chemical Society* **2017**, *140*, 362–368.
- (2) Franco, A. A.; Rucci, A.; Brandell, D.; Frayret, C.; Gaberscek, M.; Jankowski, P.; Johansson, P. Boosting Rechargeable Batteries R&D by Multiscale Modeling: Myth or Reality? *Chemical Reviews* **2019**, *119*, 4569–4627.
- (3) Yang, L.; Horton, J. T.; Payne, M. C.; Penfold, T. J.; Cole, D. J. Modeling Molecular Emitters in Organic Light-Emitting Diodes with the Quantum Mechanical Bespoke Force Field. *Journal of Chemical Theory and Computation* **2021**, *17*, 5021–5033.
- (4) Olivier, Y.; Sancho-Garcia, J.-C.; Muccioli, L.; D’Avino, G.; Beljonne, D. Computational Design of Thermally Activated Delayed Fluorescence Materials: The Challenges Ahead. *The Journal of Physical Chemistry Letters* **2018**, *9*, 6149–6163.
- (5) Nyman, J.; Pundyke, O. S.; Day, G. M. Accurate force fields and methods for modelling organic molecular crystals at finite temperatures. *Physical Chemistry Chemical Physics* **2016**, *18*, 15828–15837.
- (6) Huggins, D. J.; Biggin, P. C.; Dämgen, M. A.; Essex, J. W.; Harris, S. A.; Henchman, R. H.; Khalid, S.; Kuzmanic, A.; Laughton, C. A.; Michel, J.; Mulholland, A. J.; Rosta, E.; Sansom, M. S. P.; van der Kamp, M. W. Biomolecular simulations: From dynamics and mechanisms to computational assays of biological activity. *WIREs Computational Molecular Science* **2018**, *9*.
- (7) Mey, A. S.; Allen, B. K.; Macdonald, H. E. B.; Chodera, J. D.; Hahn, D. F.; Kuhn, M.; Michel, J.; Mobley, D. L.; Naden, L. N.; Prasad, S.; Rizzi, A.; Scheen, J.; Shirts, M. R.;

- Tresadern, G.; Xu, H. Best Practices for Alchemical Free Energy Calculations [Article v1.0]. *Living Journal of Computational Molecular Science* **2020**, *2*.
- (8) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and testing of a general amber force field. *Journal of Computational Chemistry* **2004**, *25*, 1157–1174.
- (9) Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, I.; MacKerell Jr., A. CHARMM General Force Field (CGenFF): A force field for drug-like molecules compatible with the CHARMM all-atom additive biological force fields. *J. Comp. Chem.* **2010**, *31*, 671–690.
- (10) Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *Journal of the American Chemical Society* **1996**, *118*, 11225–11236.
- (11) Qiu, Y. et al. Development and Benchmarking of Open Force Field v1.0.0—the Parsley Small-Molecule Force Field. *Journal of Chemical Theory and Computation* **2021**, *17*, 6262–6280.
- (12) Boulanger, E.; Huang, L.; Rupakheti, C.; MacKerell, A. D.; Roux, B. Optimized Lennard-Jones Parameters for Druglike Small Molecules. *Journal of Chemical Theory and Computation* **2018**, *14*, 3121–3131.
- (13) Yin, J.; Henriksen, N. M.; Slochower, D. R.; Shirts, M. R.; Chiu, M. W.; Mobley, D. L.; Gilson, M. K. Overview of the SAMPL5 host–guest challenge: Are we doing better? *Journal of Computer-Aided Molecular Design* **2016**, *31*, 1–19.
- (14) Xu, P.; Guidez, E. B.; Bertoni, C.; Gordon, M. S. Perspective: Ab initio force field methods derived from quantum mechanics. *The Journal of Chemical Physics* **2018**, *148*, 090901.

- (15) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chemical Reviews* **2021**, *121*, 10142–10186.
- (16) Barone, V.; Cacelli, I.; Mitri, N. D.; Licari, D.; Monti, S.; Prampolini, G. Joyce and Ulysses: integrated and user-friendly tools for the parameterization of intramolecular force fields from quantum mechanical data. *Physical Chemistry Chemical Physics* **2013**, *15*, 3736.
- (17) Morado, J.; Mortenson, P. N.; Verdonk, M. L.; Ward, R. A.; Essex, J. W.; Skylaris, C.-K. ParaMol: A Package for Automatic Parameterization of Molecular Mechanics Force Fields. *Journal of Chemical Information and Modeling* **2021**, *61*, 2026–2047.
- (18) Seo, B.; Lin, Z.-Y.; Zhao, Q.; Webb, M. A.; Savoie, B. M. Topology Automated Force-Field Interactions (TAFFI): A Framework for Developing Transferable Force Fields. *Journal of Chemical Information and Modeling* **2021**, *61*, 5013–5027.
- (19) Jakalian, A.; Jack, D. B.; Bayly, C. I. Fast, efficient generation of high-quality atomic charges. AM1-BCC model: II. Parameterization and validation. *Journal of Computational Chemistry* **2002**, *23*, 1623–1641.
- (20) Dodda, L. S.; Vilseck, J. Z.; Tirado-Rives, J.; Jorgensen, W. L. 1.14*CM1A-LBCC: Localized Bond-Charge Corrected CM1A Charges for Condensed-Phase Simulations. *J. Phys. Chem. B* **2017**, *121*, 3864–3870.
- (21) Bayly, C. I.; Cieplak, P.; Cornell, W.; Kollman, P. A. A well-behaved electrostatic potential based method using charge restraints for deriving atomic charges: the RESP model. *The Journal of Physical Chemistry* **1993**, *97*, 10269–10280.
- (22) Schauperl, M.; Nerenberg, P. S.; Jang, H.; Wang, L.-P.; Bayly, C. I.; Mobley, D. L.; Gilson, M. K. Non-bonded force field model with advanced restrained electrostatic potential charges (RESP2). *Communications chemistry* **2020**, *3*, 1–11.

- (23) Verstraelen, T.; Vandenbrande, S.; Heidar-Zadeh, F.; Vanduyfhuys, L.; Van Speybroeck, V.; Waroquier, M.; Ayers, P. W. Minimal basis iterative stockholder: atoms in molecules for force-field development. *Journal of Chemical Theory and Computation* **2016**, *12*, 3894–3912.
- (24) Manz, T. A.; Limas, N. G. Introducing DDEC6 atomic population analysis: part 1. Charge partitioning theory and methodology. *RSC Advances* **2016**, *6*, 47771–47801.
- (25) Limas, N. G.; Manz, T. A. Introducing DDEC6 atomic population analysis: part 2. Computed results for a wide range of periodic and nonperiodic materials. *RSC Advances* **2016**, *6*, 45727–45747.
- (26) Riquelme, M.; Lara, A.; Mobley, D. L.; Verstraelen, T.; Matamala, A. R.; Vöhringer-Martinez, E. Hydration Free Energies in the FreeSolv Database Calculated with Polarized Iterative Hirshfeld Charges. *Journal of Chemical Information and Modeling* **2018**, *58*, 1779–1797.
- (27) Tkatchenko, A.; Scheffler, M. Accurate molecular van der Waals interactions from ground-state electron density and free-atom reference data. *Phys. Rev. Lett.* **2009**, *102*, 6–9.
- (28) Cole, D. J.; Vilseck, J. Z.; Tirado-Rives, J.; Payne, M. C.; Jorgensen, W. L. Biomolecular Force Field Parameterization via Atoms-in-Molecule Electron Density Partitioning. *J. Chem. Theory Comput.* **2016**, *12*, 2312–2323.
- (29) Horton, J. T.; Allen, A. E.; Dodda, L. S.; Cole, D. J. QUBEKit: Automating the Derivation of Force Field Parameters from Quantum Mechanics. *J. Chem. Inf. Model.* **2019**, *59*, 1366–1381.
- (30) Mohebifar, M.; Johnson, E. R.; Rowley, C. N. Evaluating Force-Field London Dispersion Coefficients Using the Exchange-Hole Dipole Moment Model. *Journal of Chemical Theory and Computation* **2017**, *13*, 6146–6157.

- (31) Manz, T. A.; Chen, T. New scaling relations to compute atom-in-material polarizabilities and dispersion coefficients: Part 2. Linear-scaling computational algorithms and parallelization. *RSC Adv.* **2019**, *9*, 33310–33336.
- (32) Chen, T.; Manz, T. A. A collection of forcefield precursors for metal–organic frameworks. *RSC Advances* **2019**, *9*, 36492–36507.
- (33) Visscher, K. M.; Geerke, D. P. Deriving Force-Field Parameters from First Principles Using a Polarizable and Higher Order Dispersion Model. *J. Chem. Theory Comput.* **2019**, *15*, 1875–1883.
- (34) Visscher, K. M.; Geerke, D. P. Deriving a Polarizable Force Field for Biomolecular Building Blocks with Minimal Empirical Calibration. *The Journal of Physical Chemistry B* **2020**, *124*, 1628–1636.
- (35) Kantonen, S. M.; Muddana, H. S.; Schauperl, M.; Henriksen, N. M.; Wang, L. P.; Gilson, M. K. Data-Driven Mapping of Gas-Phase Quantum Calculations to General Force Field Lennard-Jones Parameters. *J. Chem. Theory Comput.* **2020**, *16*, 1115–1127.
- (36) Wang, L.-P.; Chen, J.; Van Voorhis, T. Systematic parametrization of polarizable force fields from quantum chemistry data. *Journal of chemical theory and computation* **2013**, *9*, 452–460.
- (37) Allen, A. E.; Payne, M. C.; Cole, D. J. Harmonic Force Constants for Molecular Mechanics Force Fields via Hessian Matrix Projection. *J. Chem. Theory Comput.* **2018**, *14*, 274–281.
- (38) Manz, T. A.; Sholl, D. S. Improved Atoms-in-Molecule Charge Partitioning Functional for Simultaneously Reproducing the Electrostatic Potential and Chemical States in Periodic and Nonperiodic Materials. *Journal of Chemical Theory and Computation* **2012**, *8*, 2844–2867.

- (39) Lee, L. P.; Limas, N. G.; Cole, D. J.; Payne, M. C.; Skylaris, C. K.; Manz, T. A. Expanding the scope of density derived electrostatic and chemical charge partitioning to thousands of atoms. *J. Chem. Theory Comput.* **2014**, *10*, 5377–5390.
- (40) Cole, D. J.; Cabeza De Vaca, I.; Jorgensen, W. L. Computation of protein-ligand binding free energies using quantum mechanical bespoke force fields. *MedChemComm* **2019**, *10*, 1116–1120.
- (41) Horton, J. T.; Allen, A. E.; Cole, D. J. Modelling flexible protein-ligand binding in p38 α MAP kinase using the QUBE force field. *ChemComm* **2020**, *56*, 932–935.
- (42) Nelson, L.; Bariami, S.; Ringrose, C.; Horton, J. T.; Kurdekar, V.; Mey, A. S. J. S.; Michel, J.; Cole, D. J. Implementation of the QUBE Force Field in SOMD for High-Throughput Alchemical Free-Energy Calculations. *Journal of Chemical Information and Modeling* **2021**, *61*, 2124–2130.
- (43) Fröhlking, T.; Bernetti, M.; Calonaci, N.; Bussi, G. Toward empirical force fields that match experimental observables. *The Journal of Chemical Physics* **2020**, *152*, 230902.
- (44) Fennell, C. J.; Wymer, K. L.; Mobley, D. L. A Fixed-Charge Model for Alcohol Polarization in the Condensed Phase, and Its Role in Small Molecule Hydration. *The Journal of Physical Chemistry B* **2014**, *118*, 6438–6446.
- (45) Boothroyd, S.; Wang, L.-P.; Mobley, D.; Chodera, J.; Shirts, M. The Open Force Field Evaluator: An automated, efficient, and scalable framework for the estimation of physical properties from molecular simulation. **2021**,
- (46) Oliveira, M. P.; Andrey, M.; Rieder, S. R.; Kern, L.; Hahn, D. F.; Riniker, S.; Horta, B. A. C.; Hünenberger, P. H. Systematic Optimization of a Fragment-Based Force Field against Experimental Pure-Liquid Properties Considering Large Compound Families: Application to Saturated Haloalkanes. *Journal of Chemical Theory and Computation* **2020**, *16*, 7525–7555.

- (47) Pierro, M. D.; Mugnai, M. L.; Elber, R. Optimizing Potentials for a Liquid Mixture: A New Force Field for a tert-Butanol and Water Solution. *The Journal of Physical Chemistry B* **2014**, *119*, 836–849.
- (48) Frisch, M. J. et al. Gaussian 09 Revision A.2. 2009.
- (49) Wang, J.; Wang, W.; Kollman, P. A.; Case, D. A. Automatic atom type and bond type perception in molecular mechanical calculations. *Journal of molecular graphics and modelling* **2006**, *25*, 247–260.
- (50) Gao, X.; Ramezanghorbani, F.; Isayev, O.; Smith, J. S.; Roitberg, A. E. TorchANI: A Free and Open Source PyTorch-Based Deep Learning Implementation of the ANI Neural Network Potentials. *Journal of Chemical Information and Modeling* **2020**, *60*, 3408–3415.
- (51) Bannwarth, C.; Caldeweyher, E.; Ehlert, S.; Hansen, A.; Pracht, P.; Seibert, J.; Spicher, S.; Grimme, S. Extended tight-binding quantum chemistry methods. *WIREs Computational Molecular Science* **2020**, *11*, e1493.
- (52) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology* **2017**, *13*, e1005659.
- (53) Smith, D. G. A. et al. PSI4 1.4: Open-source software for high-throughput quantum chemistry. *The Journal of Chemical Physics* **2020**, *152*, 184108.
- (54) Smith, D. G.; Lolinto, A. T.; Glick, Z. L.; Lee, J.; Alenaizan, A.; Barnes, T. A.; Borca, C. H.; Di Remigio, R.; Dotson, D. L.; Ehlert, S., et al. Quantum Chemistry Common Driver and Databases (QCDB) and Quantum Chemistry Engine (QCEngine): Automation and interoperability among computational chemistry programs. *The Journal of Chemical Physics* **2021**, *155*, 204801.

- (55) Mobley, D. L.; Dumont, É.; Chodera, J. D.; Dill, K. A. Comparison of Charge Models for Fixed-Charge Force Fields: Small-Molecule Hydration Free Energies in Explicit Solvent. *The Journal of Physical Chemistry B* **2007**, *111*, 2242–2254.
- (56) Prentice, J. C. A. et al. The ONETEP linear-scaling density functional theory program. *The Journal of Chemical Physics* **2020**, *152*, 174111.
- (57) Womack, J. C.; Anton, L.; Dziedzic, J.; Hasnip, P. J.; Probert, M. I. J.; Skylaris, C.-K. DL-MG: A Parallel Multigrid Poisson and Poisson–Boltzmann Solver for Electronic Structure Calculations in Vacuum and Solution. *Journal of Chemical Theory and Computation* **2018**, *14*, 1412–1432.
- (58) MacKerell, A. D. et al. All-Atom Empirical Potential for Molecular Modeling and Dynamics Studies of Proteins. *The Journal of Physical Chemistry B* **1998**, *102*, 3586–3616.
- (59) Nikitin, A. Non-zero Lennard-Jones parameters for the Toukan–Rahman water model: more accurate calculations of the solvation free energy of organic substances. *Journal of Computer-Aided Molecular Design* **2019**, *34*, 437–441.
- (60) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *The Journal of Chemical Physics* **1983**, *79*, 926–935.
- (61) Yan, X. C.; Robertson, M. J.; Tirado-Rives, J.; Jorgensen, W. L. Improved Description of Sulfur Charge Anisotropy in OPLS Force Fields: Model Development and Parameterization. *The Journal of Physical Chemistry B* **2017**, *121*, 6626–6636.
- (62) Gutiérrez, I. S.; Lin, F.-Y.; Vanommeslaeghe, K.; Lemkul, J. A.; Armacost, K. A.; Brooks, C. L.; MacKerell, A. D. Parametrization of halogen bonds in the CHARMM general force field: Improved treatment of ligand–protein interactions. *Bioorganic & Medicinal Chemistry* **2016**, *24*, 4812–4825.

- (63) Kolář, M. H.; Hobza, P. Computer Modeling of Halogen Bonds and Other σ -Hole Interactions. *Chemical Reviews* **2016**, *116*, 5155–5187.
- (64) Qiu, Y.; Smith, D. G.; Stern, C. D.; Feng, M.; Jang, H.; Wang, L. P. Driving torsion scans with wavefront propagation. *J. Chem. Phys.* **2020**, *152*, 244116.
- (65) Allen, A. E.; Robertson, M. J.; Payne, M. C.; Cole, D. J. Development and Validation of the Quantum Mechanical Bespoke Protein Force Field. *ACS Omega* **2019**, *4*, 14537–14550.
- (66) Robertson, M. J.; Tirado-Rives, J.; Jorgensen, W. L. Improved Peptide and Protein Torsional Energetics with the OPLS-AA Force Field. *J. Chem. Theory Comput.* **2015**, *11*, 3499–3509.
- (67) Řezáč, J.; Bím, D.; Gutten, O.; Rulíšek, L. Toward Accurate Conformational Energies of Smaller Peptides and Medium-Sized Macrocycles: MPCONF196 Benchmark Energy Data Set. *Journal of Chemical Theory and Computation* **2018**, *14*, 1254–1266.
- (68) Foresman, J. B.; Keith, T. A.; Wiberg, K. B.; Snoonian, J.; Frisch, M. J. Solvent Effects. 5. Influence of Cavity Shape, Truncation of Electrostatics, and Electron Correlation on ab Initio Reaction Field Calculations. *The Journal of Physical Chemistry* **1996**, *100*, 16098–16104.
- (69) Tomasi, J.; Mennucci, B.; Cammi, R. Quantum Mechanical Continuum Solvation Models. *Chemical Reviews* **2005**, *105*, 2999–3094.
- (70) Mobley, D. L.; Bannan, C. C.; Rizzi, A.; Bayly, C. I.; Chodera, J. D.; Lim, V. T.; Lim, N. M.; Beauchamp, K. A.; Slochower, D. R.; Shirts, M. R.; Gilson, M. K.; Eastman, P. K. Escaping Atom Types in Force Fields Using Direct Chemical Perception. *Journal of Chemical Theory and Computation* **2018**, *14*, 6076–6092.

- (71) Boothroyd, S.; Madin, O.; Mobley, D.; Wang, L.-P.; Chodera, J.; Shirts, M. Improving force field accuracy by training against condensed phase mixture properties. **2021**,
- (72) Gapsys, V.; Pérez-Benito, L.; Aldeghi, M.; Seeliger, D.; van Vlijmen, H.; Tresadern, G.; de Groot, B. L. Large scale relative protein ligand binding affinities using non-equilibrium alchemy. *Chemical Science* **2020**, *11*, 1140–1152.
- (73) Wang, Y.; Fass, J.; Chodera, J. D. End-to-End Differentiable Molecular Mechanics Force Field Construction. *arXiv preprint arXiv:2010.01196* **2020**,
- (74) Bleiziffer, P.; Schaller, K.; Riniker, S. Machine Learning of Partial Charges Derived from High-Quality Quantum-Mechanical Calculations. *Journal of Chemical Information and Modeling* **2018**, *58*, 579–590.