

Development and Benchmarking of Open Force Field v1.0.0—the Parsley Small-Molecule Force Field

Yudong Qiu, Daniel G. A. Smith, Simon Boothroyd, Hyesu Jang, David F. Hahn, Jeffrey Wagner, Caitlin C. Bannan, Trevor Gokey, Victoria T. Lim, Chaya D. Stern, Andrea Rizzi, Bryon Tjanaka, Gary Tresadern, Xavier Lucas, Michael R. Shirts,* Michael K. Gilson,* John D. Chodera,* Christopher I. Bayly, David L. Mobley,* and Lee-Ping Wang*



Cite This: *J. Chem. Theory Comput.* 2021, 17, 6262–6280



Read Online

ACCESS |



Metrics & More

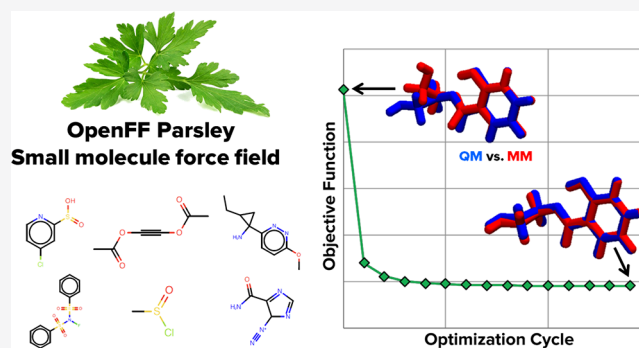


Article Recommendations



Supporting Information

ABSTRACT: We present a methodology for defining and optimizing a general force field for classical molecular simulations, and we describe its use to derive the Open Force Field 1.0.0 small-molecule force field, codenamed Parsley. Rather than using traditional atom typing, our approach is built on the SMIRKS-native Open Force Field (SMIRNOFF) parameter assignment formalism, which handles increases in the diversity and specificity of the force field definition without needlessly increasing the complexity of the specification. Parameters are optimized with the ForceBalance tool, based on reference quantum chemical data that include torsion potential energy profiles, optimized gas-phase structures, and vibrational frequencies. These quantum reference data are computed and are maintained with QCArchive, an open-source and freely available distributed computing and database software ecosystem. In this initial application of the method, we present essentially a full optimization of all valence parameters and report tests of the resulting force field against compounds and data types outside the training set. These tests show improvements in optimized geometries and conformational energetics and demonstrate that Parsley's accuracy for liquid properties is similar to that of other general force fields, as is accuracy on binding free energies. We find that this initial Parsley force field affords accuracy similar to that of other general force fields when used to calculate relative binding free energies spanning 199 protein–ligand systems. Additionally, the resulting infrastructure allows us to rapidly optimize an entirely new force field with minimal human intervention.



1. INTRODUCTION


Molecular mechanics (MM) force fields are empirical models of molecular potential energy surfaces, which yield potential energy and atomic forces as a function of atomic positions. Force fields are a crucial component of molecular simulations in many domains of chemistry and biophysics. In particular, force fields are used in simulations of biomolecular systems that may include biopolymers, aqueous solvent, and small molecules such as metabolites and therapeutics. They are also fundamental to technologies used in computer-aided drug design, such as molecular docking^{1–7} and simulation-based calculations of protein–ligand binding free energies.^{8–15}

Decades of work have led to relatively refined force fields for proteins made up of the 20 common amino acids.^{16–24} However, it is more difficult to develop a high-quality general force field, that is, one that applies to a wide range of small, organic molecules of interest in biology and drug discovery, due to the high diversity of the chemical space that must be considered. For example, nearly 100 million compounds in the PubChem database²⁵ embody many different combinations of

varied functional groups and heterocycles. Moreover, some of the most important applications of force fields involve the simulation of as-yet undiscovered compounds, such as those under investigation for small-molecule drug development, which may lie in new regions of the chemical space. Small-molecule force fields in wide use today include the general AMBER force field (GAFF),²⁶ the CHARMM general force field (CGenFF),²⁷ and the optimized potentials for liquid simulation force field (OPLS).²⁸ These important tools have undergone continuous development, and current generations are applicable to a wide range of small molecules. Nonetheless, calculations of hydration free energies, partition coefficients, and other properties show

Received: June 8, 2021

Published: September 22, 2021



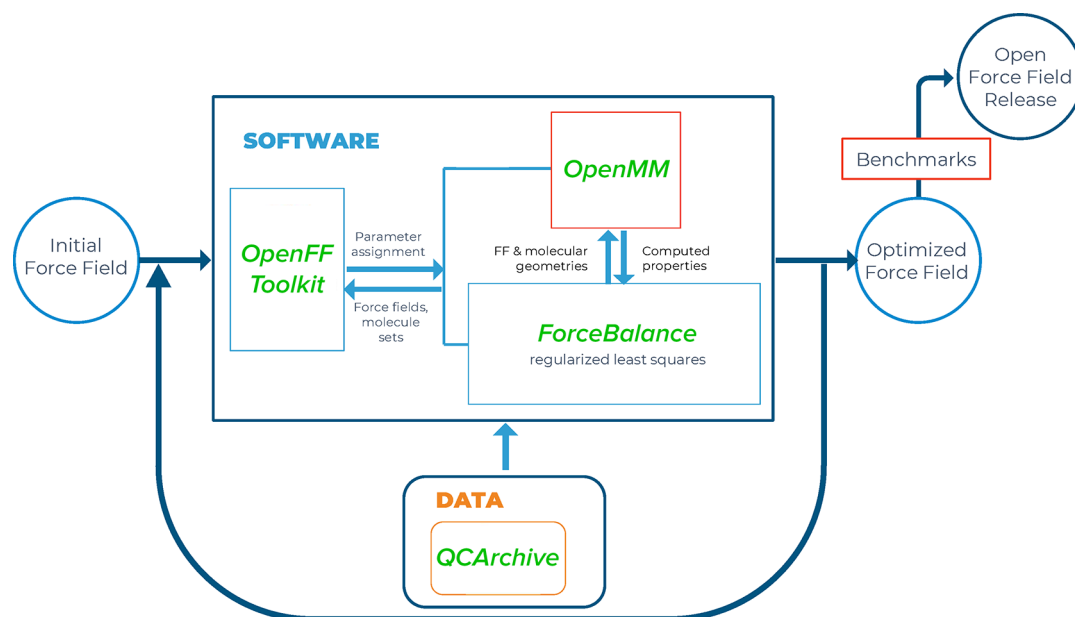


Figure 1. Open Force Field infrastructure and data flows during force field development. The OpenFF toolkit (left) sets up an MM simulation system from a given force field definition and parameters, and the OpenMM simulation code (top) is called to evaluate target physical properties. ForceBalance (middle) iteratively optimizes the parameters by least-squares minimization of an objective function constructed from the differences between MM-simulated properties and reference data. QCArchive (bottom) is a distributed computing environment and database for generation and storage of quantum chemistry reference data. See text for further details.

that there is room for improvement in current general force fields.^{29–32} In addition, weaknesses in the small-molecule component of the potential function likely account for some of the errors in calculations of noncovalent binding free energies relevant in host–guest chemistry³³ and drug design.³⁴

Improvements in force fields may come from the use of more and/or better reference data to optimize force field parameters, changes in the chemical typing rules used to assign parameters to atoms, and/or changes in the functional form of the force field. Exploring such improvements has traditionally involved considerable human input, and limitations in available computer power have made it difficult or impossible to carry out systematic explorations and optimizations. This has led to uncertainties in how exactly to go about attaining greater accuracy. For example, until it is clear how much accuracy can be achieved with a given functional form, it is impossible to ascertain whether or in what cases more accuracy is achieved by adding more detail and hence greater computational cost, for example, by accounting explicitly for electronic polarizability. In addition, it is often unclear how the specific types and parameters of a given force field were arrived at, and these conditions pose obstacles in reproducing and building on prior work. There is thus a need not only for improved force fields but also for an infrastructure that will enable systematic exploration, optimization, and evaluation of new simulation force fields. Today, advances in software technology and increasing computer power make it possible to move toward the systematization and automation of force field generation.

The present study describes a significant step in this direction—the optimization of a small-molecule force field using an automated and reproducible procedure, with all software, data, and workflows made freely available to the maximum extent possible. The result, OpenFF 1.0.0, code-named Parsley, is the first optimized force field using the SMIRNOFF format, with direct chemical perception.³⁵ To create Parsley, we started with an initial force field called

SMIRNOFF99Frosst consisting of direct chemical perception typing rules and parameters adopted from the Parm@Frosst force field.³⁶ We then optimized nearly all 500 of the valence parameters to improve agreement with quantum chemical-optimized geometries, energetics, and vibrational frequencies, while largely retaining the Lennard-Jones and electrostatic parameters of SMIRNOFF99Frosst. In keeping with the Open Force Field Initiative’s core philosophy, the Parsley force field and the software tools and data sets used to develop it are released under permissive open-source licenses.

We also report here on the initial benchmarking of Parsley to show its improved accuracy relative to its predecessor for a wide variety of properties, especially energetics and geometries relative to gas-phase quantum chemical calculations. The quantum chemical benchmarks, which cover more than 2000 molecules and probe the quality of optimized geometries and relative conformer energies, show substantially better performance relative to SMIRNOFF99Frosst. For the condensed phase properties, which span density, dielectric constant, heat of vaporization, excess molar volume, and enthalpy of mixing, no dramatic performance differences relative to SMIRNOFF99Frosst were noted in this release, and the overall accuracy is similar to that of GAFF; this was expected as nonbonded parameter optimization was not included as part of this work. In addition, as a reality check for the critical application area of computer-aided drug design, we report Parsley’s favorable performance in benchmark protein–ligand binding free energy calculations covering 199 protein–ligand pairs, along with comparative results for other small-molecule force fields. Importantly, the infrastructure described here establishes a foundation for going far beyond Parsley, through the ongoing creation of a series of continually improving, open, small-molecule force fields. The process reported here in the OpenFF 1.0.0 release continues in subsequent releases (OpenFF 1.1, 1.2, and 1.3),^{37–41} which will be described in follow-up work.




Figure 2. Illustrative subset of small fragment-like molecules from the Roche Set.

2. METHODS

The infrastructure used to generate Parsley takes an initial force field as its starting point and optimizes it against reference training data to create an optimized force field, which in turn is benchmarked against test set data prior to release (Figure 1). The software part of this infrastructure comprises a toolkit that assigns force field parameters to molecules of interest (OpenFF Toolkit); a component that computes a set of target properties for a set of input molecules (OpenMM); and the ForceBalance code, which optimizes force field parameters against the selected reference data. In general, the reference data set can include both quantum chemical and experimental data. The Parsley force field was generated by refitting the parameters in the valence terms of the initial SMIRNOFF99Frosst³⁵ force field against an extensive new set of high-level quantum mechanical data, which include energies, gas-phase geometries, and other properties. Note that the starting parameters in SMIRNOFF99Frosst in turn originated from two parent force fields, AMBER parm99⁴² and Merck's parm@Frosst,³⁶ which had been parameterized to reproduce gas-phase geometries and energetics computed at lower levels of QM than those used here, for selected molecules. Here, Section 2.1 details the force field parameters that were optimized, the QM data set used to drive the optimization, and the application of ForceBalance⁴³ to carry out the optimization from the SMIRNOFF99Frosst³⁵ starting point. Section 2.2 then describes how the resulting Parsley force field was tested against

benchmark data outside the training sets, comprising gas-phase properties from QM calculations, measured liquid-state properties, and measured protein–ligand binding free energies.

2.1. Training the Parsley Force Field. *2.1.1. Parameters that Were Refit.* There are 500 valence parameters in SMIRNOFF99Frosst, and we aimed to optimize as many of these as possible. The quantum chemical data sets that we constructed (Section 2.1.2) were successful at covering 481 of these parameters. The remaining 19 (Table S1), which were not modified due to the absence of training data (both in eMolecules, which we drew our coverage set from, and in our other available training data), either describe chemistries that are exceedingly rare in drug-like compounds or were superseded by other higher-priority parameters that matched the same chemical patterns.

All 481 parameters were fitted simultaneously against all QM data (Section 2.1.2). Each parameter definition is uniquely identified by an interaction type (e.g., bond stretch) and a SMIRKS pattern (e.g., [#6X4:1]-[#6X4:2]) and can contain one or more physical values (e.g., the bond length and the force constant). The full list of parameter definitions, which can be viewed in the published force field XML file, openff-1.0.0.offxml,⁴⁴ may be summarized as follows:

- Harmonic bond stretch: 86 equilibrium bond lengths and force constants.




Figure 3. Illustrative subset of molecules from the Coverage Set.

- Harmonic angle bend: 35 equilibrium angles and 39 force constants. These two numbers differ because four angles are linear and were kept linear during fitting.
- Proper torsions: Each of the 154 torsion types is associated with an N -term Fourier series of potential energy contributions, where $N \leq 6$, and each term, i , is of the form $E_i = k_i \cos(i\phi + \delta_i)$. We optimized all of the amplitudes that were defined in SMIRNOFF99Frosst, comprising 154, 62, 26, 5, 4, and 3 values of k_1, k_2, k_3, k_4, k_5 , and k_6 , respectively, for a total of 254 parameters. Parameters t156, t157, and t158 represent torsion angles containing a linear angle, and their values of k_1 were kept at 0.0 during fitting. The phase parameters, δ_i , and the selection of Fourier terms used for each torsion were not optimized in this release.
- Improper torsions: The four improper terms were kept unmodified, to avoid overfitting.

2.1.2. Compound Sets Used in Training. Two sets of small organic molecules were used to generate the quantum chemical data sets used in fitting. The first, termed the Roche Set, contains 468 fragment-sized molecules, most containing one to three rotatable bonds that were provided by Roche as a collection of important and/or interesting chemistries. This data set was prepared using MOE software.⁴⁵ Representative compounds from this set are illustrated in Figure 2, and the full set can be found in Supporting Information Section 1.1.2. The second, termed the Coverage Set, contains 80 molecules selected from

eMolecules⁴⁶ using a greedy algorithm aimed at providing parameter coverage for the maximum number of parameters using the minimum number of molecules. Figure 3 illustrates representative compounds, and a full list of SMILES can be found in Supporting Information Section 1.1.2.

Initial automated selection of the Coverage Set is described in a subdirectory of the openforcefield GitHub repository, and details of the additional molecules added manually to cover remaining gaps can be found in Supporting Information Section 1.1.2.

2.1.3. Selection of Quantum Chemistry Methodology. Quantum chemical calculations (geometry optimizations and torsion scans) were performed on a distributed set of high-performance computing clusters using the MolSSI QCFractal⁴⁷ distributed quantum chemistry engine, with results deposited in the public MolSSI QCArchive server (MQCAS)^{48,49} to allow open public access to all data. We used a single level of theory for all QM calculations, B3LYP-D3(BJ)/DZVP.^{50–53} This choice was based on two benchmark studies of conformational energies^{54,55} and our own initial studies that aimed to balance accuracy against computational cost. The molecules in both of these studies included amino acids, small-to-medium-sized peptides, and macrocycles. Geometries were optimized at the MP2/cc-pVTZ level, and reference energies were computed using explicitly correlated focal point analysis methods considered to be equivalent to complete basis CCSD(T) in accuracy. Both studies found that B3LYP-D3(BJ) reproduces

the reference energies with root-mean-square errors (RMSEs) of <1 kcal/mol when very large basis sets (e.g., def2-QZVP⁵⁶) are used; the empirical D3 dispersion term played a major role, as the errors were typically 2–4× larger without it.

Notably, Rezáč et al. 2018⁵⁴ reported that the double- ζ -quality DZVP basis set⁵⁰ gave nearly the same RMSE as def2-TZVP, which we were able to reproduce in our own tests. When similar-sized and better known basis sets such as 6-31G* and def2-SV(P) were used, the RMSEs increased significantly, but there were only minor differences in the computational cost. Our results are largely consistent with those of Rezáč et al. 2018,⁵⁴ even though we did not use the custom empirical dispersion parameters they derived for the DZVP-DFT basis set. A scatter plot of RMSE versus calculation time for a representative molecule, labeled as FGG114 in Rezáč et al. 2018,⁵⁴ is shown in Figures 4; the results confirm that the DZVP-DFT basis set gives the best compromise between accuracy and computational cost.




Figure 4. Tradeoff between speed and accuracy in selecting the quantum chemical basis set. Computational time (for a single conformer) vs RMSE to benchmark-quality relative energies for 15 conformations of a representative molecule for several choices of basis sets. The benchmark relative energies are MP2/CBS with a CCSD(T) correction and were obtained from ref 54.

Although we believe that our choice of the QM method is appropriate for gas-phase conformational energies for the neutral compounds comprising our training set here, we did not conduct benchmark studies on optimized geometries and vibrational frequencies, which were also part of our parameterization data set. We did not include charged molecules in the current training or test sets because our use of QM gas-phase calculations as training data does not account for any differences in properties between the gas phase and the aqueous solvent, which would be more pronounced in charged species. Therefore, the accuracy of this force field for charged species is presently unknown. This issue has been addressed in recent work on protein force fields in which the QM training data and MM calculations are both carried out in the presence of implicit solvent models,⁵⁷ which may also be applicable to small-molecule force fields. More comprehensive benchmarks are planned to inform future force field generations. However, the present level of theory is superior to the HF/6-31G* approach used in parameterizing the parm99/parm@Frosst force fields from which SMIRNOFF99Frosst descended and thus should afford greater accuracy.

2.1.4. Generation of Quantum Chemical Data for Compound Data Sets. Prior to running quantum chemical calculations, the input molecules were subjected to the

protonation state and conformer expansion, using the Fragmenter software package.⁵⁸ After the expansion, each protonation state was identified as a new molecule, so the number of distinct molecules increased; each molecule could have one or more conformers. Each conformer provided one optimized geometry used in fitting. Three classes of gas-phase quantum chemical data were generated for both the Roche and Coverage compound sets: optimized geometries, vibrational frequencies, and torsional energy profiles (Table 1). The

Table 1. Summary of Quantum Chemical Calculations Used to Fit the Force Field Valence Parameters in This Work^a

		Roche Set	Coverage Set
	compounds	468	80
	cmpds × prot. states	468	233
opt. geom.	geometries	936	831
	data set name	OpenFF Optimization Set 1	SMIRNOFF Coverage Set 1
vib. freq.	frequency sets	660	235
	data set name	OpenFF Optimization Set 1	SMIRNOFF Coverage Set 1
tors. scans	energy profiles	669	417
	data set name	OpenFF Group 1 Torsions	SMIRNOFF Coverage Torsion Set 1

^aThese publicly available data sets are stored on the MQCAS.

methods used are detailed below. The results of all quantum chemical calculations are stored as Data Set objects on MQCAS⁴⁸ and are freely available to the public. Examples of working with several MQCAS data sets can be found in Supporting Information Section 1.1.3.

2.1.4.1. Optimized Geometries. We used the MQCAS to store and distribute geometry optimizations with the geomTRIC optimization driver⁵⁹ and the Psi4 quantum chemistry package^{60,61} as backends. Optimized QM geometries were downloaded from the MQCAS and then filtered to remove cases where the bonding pattern changed on optimization and issues which pose other problems for the OpenFF toolkit v0.4.1,⁶² for example, undefined stereochemistry, missing torsion terms, or inability to assign AM1-BCC charges. Details can be found in Supporting Information Section 1.1.3.

The objective function that measures deviations of MM from QM geometries is designed in the following way: MM geometry optimizations are first locally minimized, starting from QM-optimized structures; then, MM and QM Cartesian coordinates are converted to lists of bond lengths, bond angles, and both proper and improper torsion angles. The difference between QM- and MM-optimized internal coordinates for a single molecule contributes to the objective function as

$$L_{\text{optgeo}}(\theta) = \sum_{i \in \text{ICs}} \left(\frac{x_i^{\text{MM}}(\theta) - x_i^{\text{QM}}}{d_i} \right)^2 \quad (1)$$

where θ stands for the force field parameters used in the MM calculation, d_i refers to scaling factors of 0.05 Å and 8 and 20° for bond lengths, bond angles, and improper torsion angles, respectively. Proper torsion angles were not considered here but instead were fitted based on comprehensive torsional energy profiles, as detailed below.

2.1.4.2. Vibrational Frequencies. For each optimized geometry in the Roche and Coverage molecule sets, Hessian calculations were both executed and stored in the MQCAS. From the calculations that were completed, the Hessians for the lowest energy conformation of each compound/protonation state were retained. After screening the data set for topology changes and other errors, normal-mode analysis was performed to obtain harmonic vibrational frequencies and Cartesian displacements for the internal degrees of freedom. Details can be found in [Supporting Information Section 1.1.3](#).

The corresponding force field Hessians were computed by locally minimizing the QM geometries with the force field, followed by evaluating forces with numerical displacements (0.001 Å). Normal-mode analysis was carried out, and the QM and FF frequencies were sorted from the lowest to the highest to yield the sorted sequences $\nu_{\text{QM},i}$ and $\nu_{\text{FF},i}$ respectively. The objective function contribution for each set of normal modes was computed as the sum of squared differences of the corresponding frequencies, scaled by a factor of $d_{\text{vib}} = 200 \text{ cm}^{-1}$, as

$$L_{\text{vib}}(\theta) = \sum_i \left(\frac{\nu_{\text{QM},i} - \nu_{\text{FF},i}}{d_{\text{vib}}} \right)^2 \quad (2)$$

Because the objective function sorts the QM and MM vibrational frequencies in a simple ascending order, it does not account for possible differences between the QM and FF normal modes, that is, the eigenvectors. This approach was taken primarily for computational efficiency, but it carries the risk of creating mismatches between the MM and QM vibrational modes. To test for this possibility, we also carried out single-point comparisons of vibrational frequencies, where the MM frequencies and normal modes are permuted to maximize the overlap with QM normal modes. The permutation was carried out by computing the mass-weighted overlap matrix between each pair of QM and MM normal modes followed by solving the linear sum assignment problem. The objective function was then computed using the resulting matched pairs of QM and MM frequencies.

2.1.5. Torsional Potentials. Quantum mechanical energy profiles were generated for dihedral angles in the Roche and Coverage sets (see, for example, [Figure 5](#)). All calculations were carried out on the MQCAS, which employs TorsionDrive software to compute each torsion energy profile using a wavefront propagation procedure,⁶³ described briefly here. Multiple initial geometries were generated for each molecule via Fragmenter and provided as input at the start of the procedure. Each input structure was energy-minimized with the selected torsion angle constrained to values on a 15° resolution grid, with QCArchive managing parallel job execution and individual constrained optimizations handled by geomeTRIC/Psi4 as described above. At the conclusion of the constrained minimizations, the lowest energy structure at each grid point was used to initiate new constrained minimizations at neighboring grid points. This cycle was repeated until the grid was fully populated with constrained minimization results and no new lowest energy structures were found. In order to avoid pathologies such as bond breaking that may occur when driving torsions into sterically hindered or high-energy regions, an upper energy limit was applied such that no constrained minimizations were started from structures with energies greater than 0.05 hartrees (31.3 kcal/mol) above the minimum.




Figure 5. Example torsions selected for 1D torsion scans in the Roche TorsionDrive data set. H, white; C, gray; N, blue; O, red; S, yellow; F, green (lower middle); and Cl, green (lower left).

The set of lowest energy-constrained minimized structures for each grid point was downloaded from the MQCAS and checked for bonding topology changes; calculations that contained such changes were discarded. In addition, any scans that included a frame with an internal hydrogen bond were discarded to avoid having strong intramolecular nonbonded interactions in the gas-phase QM energy, as they would lower the accuracy of the fitted parameters for condensed phase simulations. Our work also suggests that including conformations with intramolecular hydrogen bonds can adversely affect torsion fitting by conflating the effects of strong internal electrostatic interactions (which might not be well described by point charge electrostatics, especially without simultaneous refitting) into fitted torsions. In future work, hydrogen-bonded dimers may be included along with intramolecularly hydrogen-bonded conformers in order to assure a suitable balance, which implies a need to include fitting the nonbonded and valence parameters together.⁶⁴ Hydrogen bonds were detected using the Baker Hubbard method [$\text{angle}(\text{D}-\text{H}\cdots\text{A}) > 120^\circ$ and $\text{dist}(\text{H}\cdots\text{A}) < 2.5 \text{ \AA}$], as implemented in the MDTraj package.⁶⁵ Details can be found in [Supporting Information Section 1.1.3](#).

For compounds in the Roche Set, torsional scans were generated for the 819 dihedral angles, matching all of the following conditions:

1. The center bond is not part of a ring.
2. There is at least one non-hydrogen atom on both sides of the bond.
3. Neither of the two angles involved is close to linear ($\geq 165^\circ$).

Among all torsions sharing the same center bond, the one with the largest side groups, by the number of atoms, was picked. For the compounds in the Coverage Set, we used the SMIRNOFF force field to label the torsions in each molecule and selected the first five dihedral angles that matched each torsion term for scanning. (Note, however, that the force field term t155b was added after this data set was created, so no torsion was selected for that term.)

The objective function contribution was computed as a weighted sum of squared differences between QM and MM energies. During the fitting process, each structure along the QM torsional profile was partially relaxed using the empirical force

field being optimized. These MM optimizations were started from the QM-constrained optimized structure, the four atoms defining the torsion were fixed at the QM coordinates, and all other atoms were held near the QM coordinates by applying harmonic energy restraints with a force constant of 1 kcal/mol/Å². These added harmonic restraints avoid the possibility of large structural changes in the MM structures away from the QM structures, which could make the torsional profile differences less meaningful. The restraint term is applied only in the MM energy minimization, and the MM energy at the energy-minimized geometry without the restraint term is used in subsequent steps. The QM and MM energies being compared were calculated as

$$\begin{aligned} E_{\text{QM}}(\mathbf{x}_i) &= E'_{\text{QM}}(\mathbf{x}_i) - E'_{\text{QM}}(\mathbf{x}_0) \\ E_{\text{MM}}(\mathbf{x}_i; \theta) &= E'_{\text{MM}}(\text{OptMM}(\mathbf{x}_i; \theta)) \\ &\quad - E'_{\text{MM}}(\text{OptMM}(\mathbf{x}_0; \theta)) \end{aligned} \quad (3)$$

where the primes indicate absolute energies, subscripts indicate grid point indices, \mathbf{x}_0 is the lowest energy energy-minimized structure, θ represents the MM force field parameters, and $\text{OptMM}(\mathbf{x}; \theta)$ denotes the MM-constrained optimization procedure described above. The objective function is then calculated as

$$L(\theta) = \frac{1}{d_E^2} \frac{\sum_{i \in N(\text{gridpoints})} w(E_{\text{QM}}(\mathbf{x}_i)) (E_{\text{QM}}(\mathbf{x}_i) - E_{\text{MM}}(\mathbf{x}_i; \theta))^2}{\sum_{i \in N(\text{gridpoints})} w(E_{\text{QM}}(\mathbf{x}_i))} \quad (4)$$

where $d_E = 1.0$ kcal/mol is a scaling factor. The weights are calculated by a formula that uses two cutoffs, where the weight is constant until the first cutoff (1.0 kcal/mol) and then starts to decrease, followed by a hard second cutoff at 5.0 kcal/mol, above which the weights are zero.

$$w(E) = \begin{cases} 1 & E < 1.0 \text{ kcal/mol} \\ (1 + (E - 1)^2)^{-1/2} & 1.0 \leq E < 5.0 \text{ kcal/mol} \\ 0 & E \geq 5.0 \text{ kcal/mol} \end{cases} \quad (5)$$

Our decisions regarding the fitting of torsion energy profiles warrant some discussion and comparison to published studies. Previously published force fields have either employed QM torsion scans similar to this work⁶⁶ or used custom MM simulations to generate input conformations for QM calculations.⁶⁷ During the development of a nucleic acid force field, QM and MM minimizations were carried out prior to comparing energies in a similar procedure to this work.⁶⁸ The weighting scheme we used is comparable with a previous study⁶⁹ that employed a Boltzmann distribution with $T = 2000$ K ($k_B T \approx 4.0$ kcal/mol) to assign weights in fitting torsion energy profiles.

The variety of torsion-fitting procedures used in prior studies shows that a standard procedure for carrying out this step is currently lacking. This is in part because each approach involves making a different set of compromises when fitting the approximate potential, and it is challenging to assess the impact of an approach on the accuracy of simulated properties. This and other challenges motivate the creation of automated benchmarking tools that are the subject of current research but beyond the scope of this paper.

2.1.6. Optimization Algorithm and Convergence Criteria. The parameter optimization was carried out with ForceBalance,⁴³ a Python toolkit to optimize force fields in a

systematic, reproducible, scalable, and flexible manner.^{22,43} We employed a development version of ForceBalance based on v1.6.0⁷⁰ to minimize the objective function. Support of the OpenFF force field was enabled using an interface with the OpenFF toolkit v0.4.1.⁶² The commercial OpenEye toolkit, version 2019.4.2, was used to generate .mol2 files, which are needed by ForceBalance to set up OpenFF simulations using the toolkit.

Numerical derivatives of the objective function with respect to parameters were computed with dimensionless displacements of 0.01 for improved numerical stability, relative to the ForceBalance default of 0.001. Fitting was terminated once two convergence criteria were met:

1. The dimensionless parameter step size fell below 0.01.
2. The objective function (Section 2.1.7) decreased by less than 0.1 during the step.

To efficiently optimize the parameters in as few iterations as possible, ForceBalance uses a quasi-Newton iteration to take near-optimal steps in the parameter space

$$\theta^{(n+1)} = \theta^{(n)} + [\mathbf{H}(\theta) + \lambda \mathbf{I}]^{-1}$$

To approximate the Hessian $\mathbf{H}(\theta)$, ForceBalance computes an approximation to the matrix of second derivatives of each least-squares component in a manner that neglects parameter couplings

$$\begin{aligned} \mathbf{H}_{i,pq}(\theta) &= \frac{\partial^2}{\partial \theta_p \partial \theta_q} L_i(\theta) = \sum_{j \in \text{data}} \frac{\partial A_j}{\partial \theta_p} \frac{\partial A_j}{\partial \theta_q} + \frac{\partial^2 A_j}{\partial \theta_p \partial \theta_q} \\ &\approx \sum_{j \in \text{data}} \frac{\partial A_j}{\partial \theta_p} \frac{\partial A_j}{\partial \theta_q} \end{aligned} \quad (6)$$

The λ parameter is used to set the optimization step size, and it was determined by line-search minimization for a given parametric gradient and Hessian. This strategy was employed because the line search over λ only requires repeated evaluation of the objective function itself and not the parametric gradient, which is relatively expensive.

2.1.7. Objective Function with Regularization. ForceBalance was used to minimize an objective or loss function, $L(\theta)$, with respect to force field parameters θ . The objective function quantifies deviation of quantities derived with the force field from the reference quantum chemical data, while adding a regularization penalty to minimize the deviation from a reference set of parameters, following the standard approach for ForceBalance⁴³

$$L_{\text{tot}}(\theta) = \sum_{i \in \text{targets}} w_i L_i(\theta) + w_{\text{reg}} \sum_{p \in \text{parameters}} \frac{|\Delta \theta_p|^2}{\sigma_p^2}$$

where w_i is the weight of each class of optimization data targets with the corresponding loss functions $L_i(\theta)$, which are often least-squares penalized loss

$$L_i(\theta) = \sum_{j \in \text{data}} (A_j^{\text{obs}} - A_j^{\text{calc}}(\theta))^2$$

where A_j^{obs} is an observed quantum chemical or physical property target to fit and $A_j^{\text{calc}}(\theta)$ is the calculated value. w_i of each target type was chosen to prevent the optimizer prioritizing one target type over other types. Carefully selected weights enabled the objective function contributions of different target types to have the same order of magnitude at the start of

parameter optimization. w_{reg} is the regularization penalty weight, and $\Delta\theta$ quantifies the deviation from a reference set of parameters—here, the initial SMIRNOFF99Frosst v 1.1.0 parameter set.³⁵ Regularization ensures that parameter adjustments are made conservatively to avoid introducing large problematic parameter changes that may only provide marginal improvements in the optimization target, especially when smaller data sets are used in parameterization. We used the regularization scales, σ_p , listed in Table 2, based on past observations of variations in these parameter types in previous studies.²²

Table 2. Regularization Scaling Parameters Used in ForceBalance Optimization Runs for Each Force Field Parameter Type

parameter	regularization scale σ_p
bond force constant K_c	100 kcal/mol/Å ²
bond equilibrium length r_0	0.1 Å
angle force constant K_θ	100 kcal/mol rad ²
equilibrium angle θ_0	20°
proper torsion barrier height K	1 kcal/mol
vdW well depth ϵ	0.1 kcal/mol
vdW minimum $r_{\text{min-half}}$	1 Å

2.2. Testing the Parsley Force Field. Once the parameters had been trained as detailed in Section 2.1, we tested the resulting force field, Parsley 1.0.0, against optimized geometries outside the training set and compared the results to those obtained with the initial force field, SMIRNOFF99Frosst v1.1.0.³⁵ We also tested Parsley against two data types outside the training set: energy differences among conformers of a given molecule and physical properties of various organic liquids. Tests against vibrational spectra and torsional energy potentials are reserved for future studies. Benchmark comparisons of Parsley in the context of other general force fields are also available.⁷¹ We now describe how these tests were done.

2.2.1. Quantum Chemical Test Set Generation. The QCArchive tool was used to generate and archive additional QM data, using the procedures detailed in Section 2.1, for compounds in three collections.

2.2.1.1. OpenFF Discrepancy Benchmark 1. This comprises 2802 fragment-like molecules (19,712 conformers) selected from the eMolecules catalog⁴⁶ because their energy-minimized geometries are substantially different in SMIRNOFF99Frosst 1.0.8 relative to GAFF, GAFF2, MMFF94, and MMFF94s.^{72,73} We retained all protonation and tautomer states present in our initial data set but did not generate any additional ones. Further details can be found in Supporting Information Section 1.2.1.

2.2.1.2. Pfizer Discrepancy Optimization Data Set 1. This comprises 100 fragment-like molecules, for which Pfizer's QM calculations of torsional profiles computed with HF/minix⁷⁴ followed by B3LYP/6-31G**//B3LYP/6-31G** differed substantially from those generated with the OPLS3e force field. The Pfizer code for relevant related work is public on GitHub.⁷⁵ Enumeration of conformers, but not of protonation states, led to 352 conformers.

2.2.1.3. FDA Optimization Data Set 1. This is a subset of the list of FDA-approved drugs in the ZINC15 FDA data set.⁷⁶ Molecules were retained if they had 3–55 heavy atoms and consisted only of the elements H, C, N, O, F, P, S, Cl, Br, I, and B. We retained multiple protonation and tautomer states in the database but did not generate any additional ones. Generation of

up to 20 conformers per molecule led to 6675 conformers for the 1038 structures.

Test results are presented for the merger of these three data sets, termed the Full Benchmark Set. This data set can be retrieved from the MQCAS as OpenFF Full Optimization Benchmark 1, as documented elsewhere.⁷⁷ This is an “optimization data set” in the sense that it and the results presented here for benchmarking on this set are for performance on optimized geometries only.

Conformational energy differences were assessed as follows. Compound conformers were energy-minimized using QM. For a compound with at least three conformers, we identified the conformer i_{min} with the lowest QM energy $E_{i_{\text{min}},\text{QM}}$ and computed the energies of its other conformers relative to it: $E_{i,\text{QM}} - E_{i_{\text{min}},\text{QM}}$. We then computed the force field energies of the same conformers, $E_{i,\text{FF}}$, and, for each compound, we computed the RMSE of $E_{i,\text{FF}} - E_{i_{\text{min}},\text{FF}}$ from the corresponding QM energies. Note that conformation i_{min} is based on the QM energies and used again for the FF energies.

2.2.2. Testing against Physical Properties of Organic Liquids. We assessed the ability of molecular dynamics simulations using the newly fitted Parsley force field to replicate 221 experimental observables for organic liquids spanning 104 molecules. The observables used are densities, heats of vaporization, and static dielectric constants of pure liquids and excess molar volumes and heats of mixing of binary liquid mixtures. The experimental data were drawn from the NIST ThermoML archive.⁷⁸ For systems involving water, the TIP3P model⁷⁹ was used. Automated scripts used to select the data can be found in Supporting Information Section 1.2.2.

We started with all available measurements of the properties listed above. When multiple values were available for a given quantity, only the ones with the lowest reported uncertainties were retained. We furthermore excluded ionic liquids, compounds containing elements other than H, N, C, O, S, F, Cl, Br, and I, and data measured outside the temperature range of 288–318 K and the pressure range of 0.95–1.05 atm. Dielectric constants <10 were also excluded because a force field that does not explicitly include electronic polarizability is not expected to replicate lower dielectric constants well.⁸⁰ Finally, a greedy search was performed on the remaining data to select a minimal subset of small molecules that exercised the largest number of nonbonded parameter types and for which most measurements were available. Sample compounds from the resulting set are shown in Figure 6, and further information on the data set can be found in Supporting Information Section 1.2.2.

Values for all of these properties were computed with the OpenFF-Evaluator (formerly named PropertyEstimator) 0.0.5 tool, using scripts which can be found in Supporting Information Section 1.2.2. Calculations were carried out with the new Parsley 1.0.0 tool and, for comparison, with its precursor, SMIRNOFF99frosst 1.1.0, and also GAFF 1.81 and GAFF 2.11.

2.2.3. Testing against Protein–Ligand Binding Free Energies. We assessed the performance of the newly fitted Parsley force field in binding free energy calculations based on molecular dynamics simulations following suggested best practices for benchmarking binding affinities.⁸¹ The test set consisted of eight protein targets with a total of 199 ligands (Supporting Information, Table S2), using a set commonly referred to as the “JACS data set” from a prior study published in that journal and frequently used by the community.¹² For a fair




Figure 6. Representative molecules in the condensed phase physical property benchmark set.

comparison to previously published results, the initial ligand and protein structures were obtained from prior benchmark work.⁸² These structures are available in the protein–ligand benchmark repository.⁸³ Relative binding free energies were calculated employing alchemical perturbations between pairs of ligands in water and the protein complex. These calculations employed a nonequilibrium workflow based on GROMACS and pmx.^{84,85} For the ligand molecules, the Parsley force field was used for the parameters. The protein was parameterized with the AMBER

ff99sb*ILDN force field,^{20,86,87} and a TIP3P explicit water model was employed. We chose AMBER ff99sb*ILDN as the protein force field because we expect Parsley to have the best compatibility with the AMBER family of force fields, as its nonbonded parameters are ultimately taken from Parm@Frosst, an AMBER-compatible small-molecule force field.³⁶ The water model was chosen as TIP3P due to the widespread use of this water model with the AMBER family of protein force fields. To mimic physiological conditions, ions (150 mM NaCl) and

additional counterions to neutralize the system were added to the dodecahedral simulation boxes.

The analysis workflow used for analyzing the calculations is available in ref 88. The statistics in this workflow were calculated using Arsenic,⁸⁹ which is a package implementing best practices for consistently calculating statistics and reporting results from binding free energy calculations. The Parsley results were compared to previously published results using the GAFF2.1,²⁶ CGenFF3.0.1,²⁷ and OPLS3e^{34,66} force fields. The former GAFF2.1 and CGenFF results were calculated with the same pmx workflow, and the OPLS3e results were calculated with Schrödinger FEP+.⁸²

A more detailed discussion of the workflow, the employed parameters, and the analysis can be found in [Supporting Information](#). Full details of this study will be reported in a forthcoming publication, and a preliminary report can be found in ref 90.

3. RESULTS AND DISCUSSION

This section first describes the consequences of parameter optimization for accuracy over the training set and then benchmarks Parsley on the separate test set compounds and properties. The test set results should be indicative of Parsley's accuracy in new applications.

3.1. Improvements in Accuracy over Training Set Data. **3.1.1. Optimization of the Objective Function.** The parameter fitting process dramatically increased the accuracy of the force field for the training data. Although this was anticipated, it is important to confirm because it verifies the effectiveness of the optimization procedures and provides a scale for the degree of improvement. The dimensionless objective (or loss) function—the weighted sum of squared differences between QM and MM values—decreased in the fitting from 25,708 to 4522 ([Figure 7](#)). As described in [Section 2.1](#), the




Figure 7. Objective function or loss function plotted against the number of ForceBalance iterations.

objective function is a sum of contributions, which reports the accuracy of optimized geometries, vibrational spectra, and torsional energy profiles. The effect of training on these components is summarized in [Table 3](#) (training set data) and [Figure 8](#), and the following subsections provide further details of these results. Full fitting details and inputs and outputs can be found in the release package.⁹¹

3.1.2. Optimized Geometries. The geometric component of the objective function is computed from the deviations of bond lengths, bond angles, and improper torsions, in structures optimized with the force field, from their values in the corresponding structures optimized with QM ([Section 2.1.4](#)). As shown in [Figure 8a](#), the fitting process led to an improved overall agreement between force field and QM geometries; compare the initial (red) and final (green) histograms. The portion of the blue histogram on the negative/positive x -axis shows the percentage of targets where accuracy is improved/degraded, respectively. The accuracy was somewhat reduced for a small minority of conformers, as evident from the histogram of differences (blue), but this is as expected because compromises have to be made for some molecules in order to improve the accuracy for others that use the same parameters. [Table 3](#) provides a physically interpretable perspective of these results, showing that the root-mean-square (rms) errors of bond lengths, bond angles, and improper torsions in the optimized geometries decreased by 14–49% with training.

It can also be useful to assess the fitting of individual parameters. To do this for a given bond stretch parameter, for example, we collected all test cases that included the parameter of interest and made a scatter plot of the length of the bond in the QM geometries versus the length in the MM geometries. Such plots were generated for each bond length, bond angle, and improper dihedral, in the training set, and all are available in the release package.⁹¹ When considering this term-by-term analysis, it should be kept in mind that the length of a bond or the value of an angle in an optimized geometry is determined not only by the parameters of the corresponding force field term but also by the rest of the structure. For example, a bond length may be shifted by ring strain. However, when these values consistently differ between QM and MM geometries, this can be an indicator that the specific force field parameter requires further attention.

The fitting process moved most bond lengths and angles closer to the diagonals of these scatter plots, implying better agreement between MM and QM, as expected based on [Figure 7](#). For a clear example, see [Figure 9](#) (left), where a change in the equilibrium bond length shifted the MM results to the diagonal and thus into better agreement with QM. However, a few of these scatter plots are more problematic. For example, [Figure 9](#) (right), which examines a general N=N bond stretch, shows a small shift of the data points toward the diagonal but does not correct the fact that this bond length falls in a narrow range across all the MM geometries but is varied in the QM geometries. In cases like this, greater accuracy might be achieved by creating two or more force field types for N=N bonds rather than just one. Before taking such a step, however, one should consider whether the variations in the QM bond lengths trace to varied amounts of strain placed on the bond by other components of the structure. If so, the accuracy of the N=N bond lengths should be improved by adjustment of other parameters that would correct the strain, rather than correcting parameters intrinsic to the bond itself.

Relationships among force field equilibrium bond lengths, chemical environment, and strain may be further explored by examining the lengths of a given bond type across the geometrically optimized conformers of various compounds. [Figure 10](#) illustrates this concept for a generic C–N single bond. The curves in the left panel show the energy-minimized central bond lengths in QM torsion profiles taken from the Roche data set, including all cases where the central bond is matched by the SMIRKS pattern indicated. The rise and fall of an individual

Table 3. Overall Change in RMSE Metrics vs the Quantum Chemical Result Calculated for Four Types of Properties, Using the Initial and Optimized Force field, and Divided into Training Set and Test Set^a

data class		training set			full test set		
		initial RMSE	final RMSE	change (%)	initial RMSE	final RMSE	change (%)
geometry optimization	bond lengths (Å)	0.045	0.023	-49	0.023	0.017	-33
	bond angles (deg)	3.71	3.20	-14	3.80	3.59	-5.5
	improper dihedrals (deg)	4.15	2.87	-31	4.31	2.82	-35
vibrational spectra (with mode reassignment)	frequencies (cm ⁻¹)	119 (156)	39.6 (90)	-67 (-42)	ND	ND	ND
torsion energy profiles	energies (kcal/mol)	2.96	1.89	-36	ND	ND	ND
relative energies	energies (kcal/mol)	ND	ND	ND	2.43	2.13	-12

^aThe numbers in parentheses under vibrational spectra indicate RMSE in frequencies after permutation of MMnormal modes to maximize overlap with QM normal modes. ND = No Data.

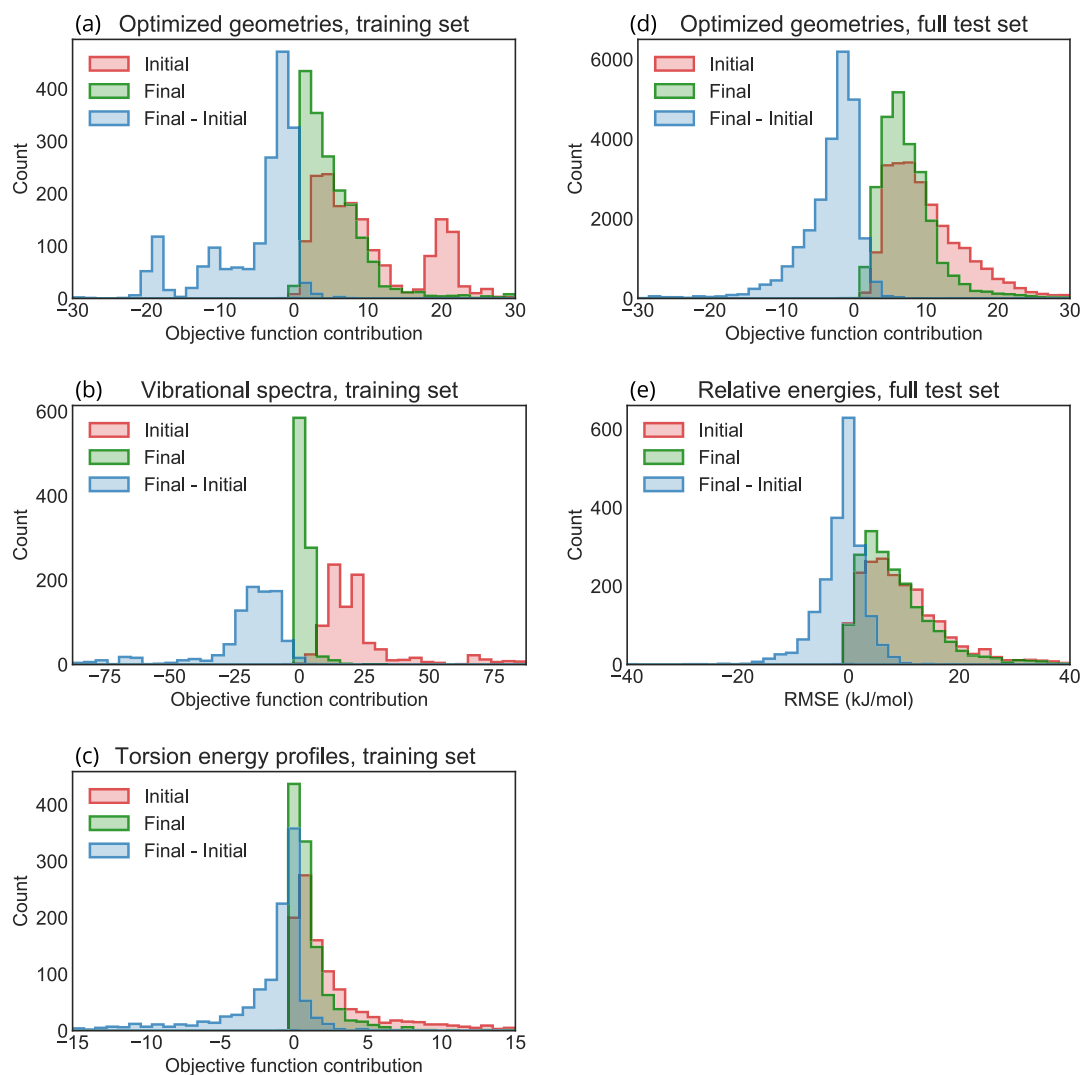


Figure 8. Improvement in components of the training set and test set objective functions with fitting. The red histogram shows performance with our initial force field, the green histogram shows performance with the optimized force field, and the blue histogram shows the distribution of changes in the objective function contribution of each target (individual molecules/geometries contributing to the objective function) due to the parameter optimization. The left column (a–c) provides the training set results. The right column (d–e) provides test set results. The range of each plot encompasses $\geq 94.94\%$ of the population of initial objective function contributions and $\geq 99.2\%$ of the population of final objective function contributions.

curve indicate the dependence of the central bond length on the torsion angle, and the vertical displacements of the various curves relative to each other indicate the torsion angle-independent differences of central bond lengths between different molecules or different central bonds in the same

molecule. When integrated over all torsion angles, the bond lengths across all instances of the central bond matched by this SMIRKS pattern have a bi- or tri-modal distribution (Figure 10, right panel). This result suggests that this generic bond type ought to be split into at least two or perhaps three more specific

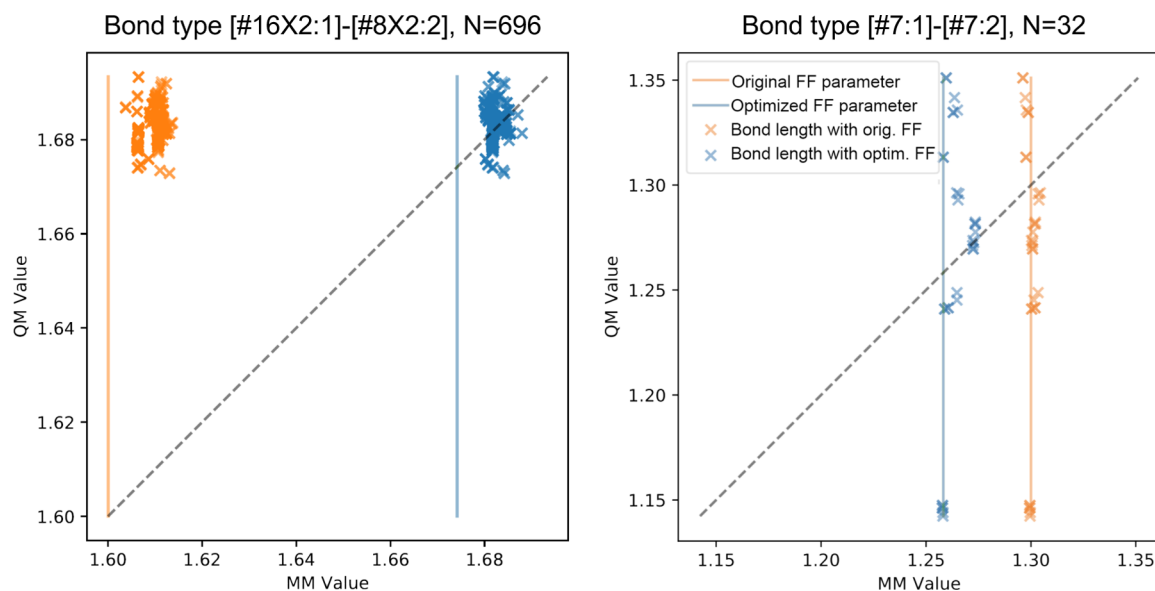


Figure 9. Comparison of QM and MM energy-minimized bond lengths for two parameters. Left: divalent sulfur single-bonded to a divalent oxygen. Right: divalent bond between nitrogens. The vertical line indicates the value of the force field's equilibrium bond length. Orange and blue indicate results for the initial and optimized force fields, respectively. Dashed: line of identity.

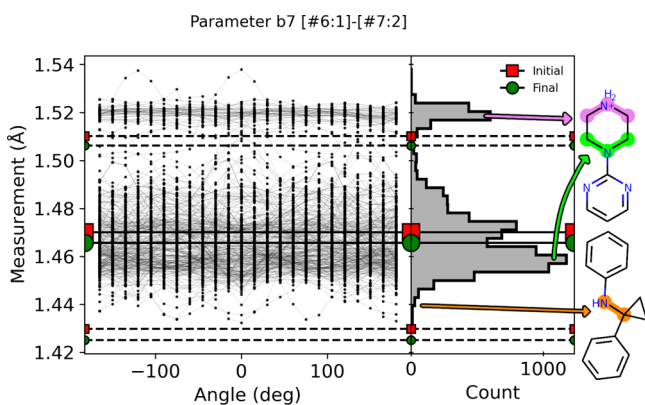


Figure 10. Dependence of bond lengths of a given force field type upon the chemical and conformational environment. Left: length of the central bond as a function of the torsion angle in the Roche data set for central bonds matching the C–N bond type indicated. Each line corresponds to the length of the C–N bond, matching the b7 parameter for constrained energy-minimized conformers over a range of torsion angles. Right: histogram of the observed bond lengths after summing over the torsion angles. The solid line in both panels indicates the force field's equilibrium bond length for type b7, and dashed lines indicate the lengths for which the bond energy equals 1.2 kcal/mol. The lines labeled in red are b7 in the initial force field, and the lines labeled in green are b7 after force field optimization. Example molecules and their given b7 bond(s) are highlighted on the far right, which correspond to typical environments where the bond length is 1.44 Å (orange), 1.46 Å (green), or 1.52 Å (violet).

types determined by SMIRKS patterns, matching more specific chemical environments.

3.1.3. Vibrational Frequencies. Fitting against the training set led to substantial improvements in the accuracy of the vibrational frequencies in the training data, relative to the reference QM results, as is evident from a dramatic fall in the vibrational components of the objective function. This is evident in Figure 8b, which shows decreases in error for the sum of squared vibrational frequency differences for individual molecules. Indeed, the improvement from initial results (red)

to fitted results (green) appears even more marked than that of optimized geometries (Figure 8a). The distribution of improvements for individual conformations (blue) also shows strong improvement, with only a tiny minority of cases becoming less accurate with fitting. These results correspond to a 67% drop in the RMSE of individual MM versus QM frequencies, that is, from 119 to 40 cm^{-1} (Table 3). When the MM vibrational frequencies are permuted to maximize the overlap between MM and QM normal modes, the RMSE in the vibrational frequencies is found to decrease by 42% (156 to 90 cm^{-1}) with optimization.

3.1.4. Torsional Energy Profiles. Fitting also led to improvements in the accuracy of the torsional energy profiles in the training set (Figure 8c), although the improvements (red to green) appear less dramatic than for the geometric and vibrational components discussed above. As in the case of the other objective function components, the improvements observed for many torsions come at the expense of decreased accuracy for some others (blue). The RMSE of the MM torsional energy profiles relative to the QM ones in the training set fell from 2.96 to 1.89 kcal/mol, a 36% drop (Table 3).

It is also of interest to compare the MM and QM potential energy profiles for individual torsion angles across the full training set, and a full set of comparisons is available in the Supporting Information. Sample plots for a torsional profile that improves with fitting and another that gets worse are provided in Figure 11, left and right panels, respectively. Interestingly, the parameter in the second plot occurs 231 times in the training set, so degraded performance is likely not due to a lack of sufficient data but instead to either changes in other portions of the force field or improved performance on other molecules utilizing this same parameter at the expense of degraded performance for this particular target. Note that most torsional parameters appear in many molecules in the training set, so fitting can improve accuracy for most occurrences, while degrading it for others.

The greater difficulty of fitting torsional profiles may result from the fact that these are particularly sensitive to nonlocal interactions within the molecules, such as longer range sterics and electrostatics. Also, defining force field types for torsional terms is more complex than for most other terms in the force

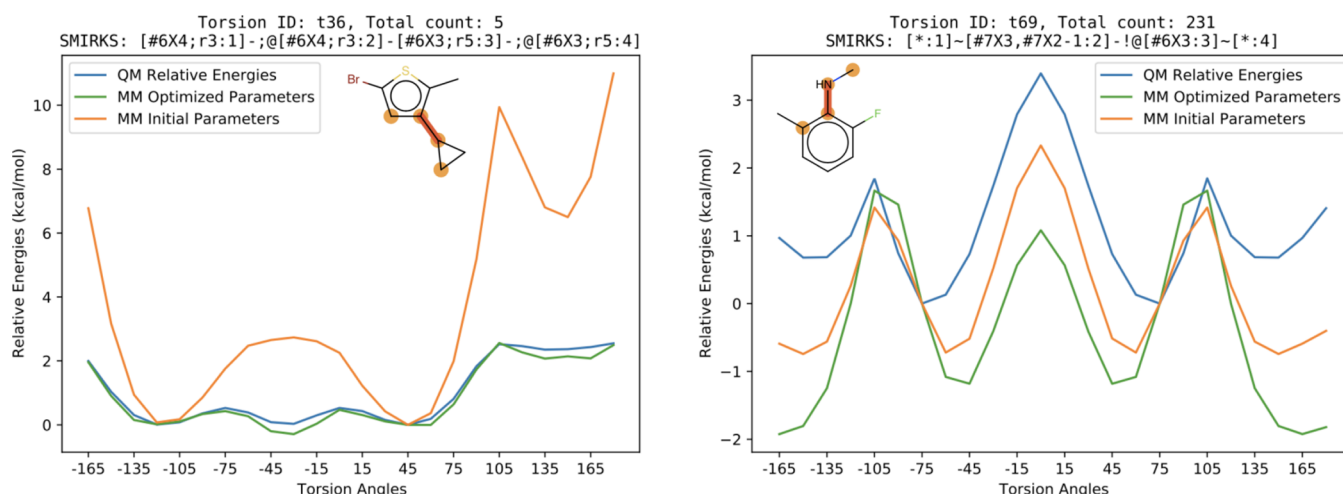


Figure 11. Examples of torsional profiles that were improved (left) or degraded (right) by fitting. Data are for a specific torsion angle in a specific molecule, as detailed below the plot. Blue: QM energy. Orange: force field energy before training. Green: force field energy after training (Parsley). The metadata at the bottom explain which data set these data are drawn from and which specific molecule this torsion occurs in and the SMIRKS pattern for the particular torsion being fitted here. The total count of this SMIRKS pattern across the data set (5) is also shown at the bottom along with the parameter ID and the atom indices in the molecule. The full set of plots is available in the release package.⁹¹

field, as multiple torsional terms contribute to the profile around a given bond, and torsional terms include step changes in periodicity. Note, also, that the present fitting process adjusted only amplitudes and left periodicities and phases unchanged. Adjustment of these additional parameters will clearly be of interest in future rounds of force field development.

3.2. Test Set Results. Results for data outside the training set provide an indication of the transferability of the new parameters and hence of the accuracy that may be expected in actual use. Here, we examine the ability of the new parameters to replicate QM-optimized gas-phase geometries for molecules outside the training set, energy differences between gas-phase conformers, physical properties of liquids, and relative protein–ligand binding free energies.

3.2.1. Quantum Chemical Data. The overall objective function for the test set is lower for Parsley (20,672) than for the initial force field (29,469). The distribution of improvements over the test set compounds (Figure 8d) shows that the objective function improves for almost all compounds, given that the blue histogram of differences has few positive values. Accordingly, improvements of 6–35% are observed in the terms that contribute to the objective function (Table 3). It is worth noting that the test set exercises 415 out of the 500 parameters. We also grouped the bond lengths, bond angles, and improper dihedrals across the test set compounds according to their FF types and examined the improvement in accuracy by the type, as illustrated for the bond lengths in Figure 12. The complete figures for bond lengths, bond angles, and improper torsion angles can be found in Supporting Information Figures S1–S3. Clearly, optimization over the training set led to improved test improvement for most parameters. Comparable plots for angles and torsions are available in the release package for this force field.⁹²

We also tested the ability of the Parsley force field to replicate differences among conformations of gas-phase molecules in the test set. Note that this type of data is entirely absent from the training data. Nonetheless, the RMSE for these quantities fell by 12% on going from the initial force field to the new Parsley force field (Table 3). The improvements in accuracy are distributed across many compounds, rather than being dominated by




Figure 12. Bond length RMSE comparison for initial and optimized force fields for the full set. For each bond type (b1, b2...), a gray circle indicates the RMSE of bonds of this type for the initial force field and arrows show the drops (green) or increases (red) in error on going to the new force field. (SMIRKS patterns for these parameter IDs can be retrieved from the force field XML file, openff-1.0.0.offxml.)⁴⁴

improvements for a few, as evident from the histograms in Figure 8e.

3.2.2. Physical Properties of Organic Liquids. We tested Parsley's ability to model condensed phase properties using it to compute densities, enthalpies of vaporization, static dielectric constants, enthalpies of mixing, and excess molar volumes of organic liquids and mixtures and comparing with experimental data from NIST's ThermoML. Note that no condensed phase data were used in the fitting process. As shown in Figures 13 and

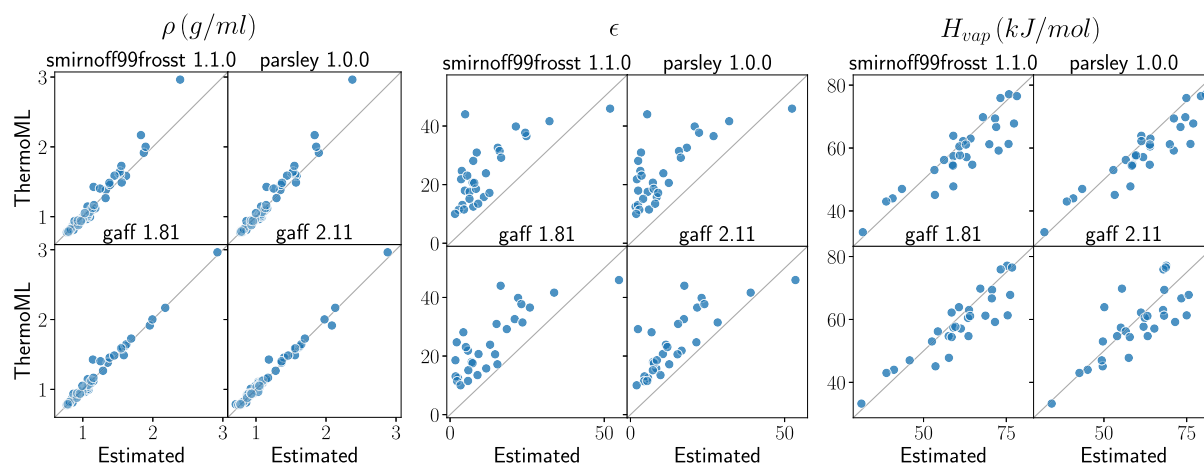


Figure 13. Results of pure property benchmarks. Liquid properties computed with various force fields, as labeled, are compared with those from the experiment. Density: ρ ; dielectric constant: ϵ ; and heat of vaporization: H_{vap} .

14, the new Parsley force field offers competitive performance for these data, with marginal, although not statistically significant (by comparison of the root-mean-square errors and their 95% confidence intervals), improvement over the previous SMIRNOFF99frosst 1.1.0 release (Table 4). The overall accuracy also

is similar to that of the established GAFF family of force fields. This pattern presumably reflects the fact that these physical properties are not sensitive to the valence parameters optimized here and that condensed phase data were not used to guide the optimization.

3.2.3. Protein–Ligand Binding Free Energies. The Parsley force field provides competitive accuracy in relative binding free energy calculations for 199 ligands across eight different protein targets, as shown in Figure 15. Indeed, the differences in accuracy across the four force fields examined here are within 95% confidence intervals. It is also important to note that the accuracy of these calculations is strongly affected by additional factors, including input structure preparation and sampling time. This said, in terms of mean unsigned error (MUE) of all relative free energy differences $\Delta\Delta G$, Parsley (MUE = 1.02 kcal/mol) ranks third, after GAFFv2.1 (MUE = 0.92 kcal/mol) and OPLS3e (MUE = 0.93 kcal/mol) and before CGenFF (MUE = 1.09 kcal/mol). These results indicate that Parsley is a reasonable choice of force field for binding free energy calculations in drug discovery projects.

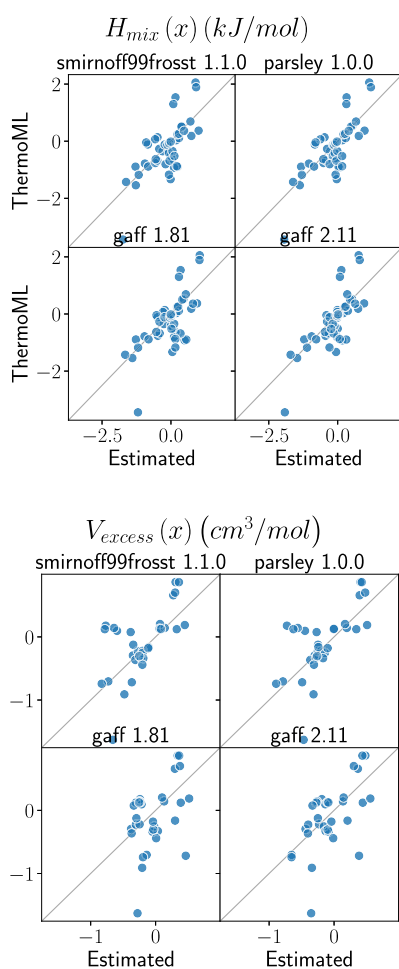


Figure 14. Results of binary mixture property benchmarks. Liquid properties computed with various force fields, as labeled, are compared with those from the experiment. Enthalpy of mixing: H_{mix} and excess molar volume of mixing: V_{excess} .

4. USING AND CITING PARSLEY

The present Parsley force field, formally named openff 1.0.0, can be accessed from Python by installing the Open Force Field toolkit with the command `conda install -c omnia openforcefield openforcefields` and then loading the force field as follows:

```
from openforcefield.typing.engines.smirnoff import ForceField
ff = ForceField('openff-1.0.0.offxml')
```

The default version of Parsley includes hydrogen bond length constraints, which allow use of the typical 2–4 fs timestep in molecular dynamics simulations. A second version without these constraints, which is suitable for geometry optimizations and single-point energy calculations, may be accessed as follows:

```
ff = ForceField('openff_unconstrained-1.0.0.offxml')
```

An example of the use of Parsley to run a molecular dynamics simulation can be found in Supporting Information Section 3. Alternatively, the force field files themselves can be found under the `openforcefields/offxml` subdirectory of the `openforcefields` GitHub repository.⁹³

The present Parsley version may be referred to as “Open Force Field (OpenFF) Parsley force field (v1.0.0)” on the first

Table 4. Measures of Accuracy of Force Fields for the Physical Property Benchmarks^a

property	force field	RMSE	R ²	τ
$V_{\text{excess}}(x)$ (cm ³ /mol)	smirnoff99frosst 1.1.0	0.41 _{0.29} ^{0.52}	0.39 _{0.14} ^{0.70}	0.44 _{0.13} ^{0.67}
	parsley 1.0.0	0.39 _{0.26} ^{0.51}	0.44 _{0.18} ^{0.71}	0.50 _{0.22} ^{0.71}
	gaff 1.81	0.47 _{0.34} ^{0.61}	0.17 _{0.01} ^{0.45}	0.23 _{-0.01} ^{0.46}
	gaff 2.11	0.41 _{0.27} ^{0.54}	0.36 _{0.10} ^{0.66}	0.42 _{0.14} ^{0.62}
$H_{\text{mix}}(x)$ (kJ/mol)	smirnoff99frosst 1.1.0	0.64 _{0.50} ^{0.76}	0.50 _{0.25} ^{0.67}	0.49 _{0.28} ^{0.66}
	parsley 1.0.0	0.58 _{0.46} ^{0.69}	0.59 _{0.34} ^{0.74}	0.54 _{0.32} ^{0.70}
	gaff 1.81	0.69 _{0.52} ^{0.87}	0.45 _{0.25} ^{0.62}	0.44 _{0.22} ^{0.63}
	gaff 2.11	0.60 _{0.45} ^{0.73}	0.58 _{0.35} ^{0.74}	0.58 _{0.40} ^{0.72}
H_{vap} (kJ/mol)	smirnoff99frosst 1.1.0	6.03 _{4.24} ^{7.56}	0.79 _{0.62} ^{0.90}	0.70 _{0.50} ^{0.83}
	parsley 1.0.0	5.87 _{4.23} ^{7.37}	0.85 _{0.71} ^{0.93}	0.79 _{0.61} ^{0.88}
	gaff 1.81	5.53 _{3.97} ^{7.12}	0.80 _{0.63} ^{0.90}	0.72 _{0.54} ^{0.84}
	gaff 2.11	7.00 _{5.33} ^{8.52}	0.60 _{0.30} ^{0.80}	0.59 _{0.35} ^{0.76}
ρ (g/mL)	smirnoff99frosst 1.1.0	0.10 _{0.06} ^{0.14}	0.96 _{0.94} ^{0.98}	0.90 _{0.85} ^{0.92}
	parsley 1.0.0	0.10 _{0.05} ^{0.15}	0.96 _{0.94} ^{0.98}	0.90 _{0.85} ^{0.92}
	gaff 1.81	0.05 _{0.03} ^{0.07}	0.98 _{0.95} ^{1.00}	0.91 _{0.87} ^{0.94}
	gaff 2.11	0.05 _{0.03} ^{0.07}	0.99 _{0.97} ^{1.00}	0.92 _{0.87} ^{0.94}
ϵ	smirnoff99frosst 1.1.0	14.78 _{11.62} ^{18.17}	0.53 _{0.22} ^{0.80}	0.50 _{0.22} ^{0.72}
	parsley 1.0.0	15.72 _{12.54} ^{19.19}	0.53 _{0.22} ^{0.79}	0.49 _{0.20} ^{0.69}
	gaff 1.81	13.22 _{10.80} ^{15.63}	0.64 _{0.44} ^{0.82}	0.54 _{0.34} ^{0.73}
	gaff 2.11	12.16 _{9.40} ^{14.96}	0.62 _{0.40} ^{0.82}	0.67 _{0.42} ^{0.81}

^aRMSE: root-mean-square error; R²: coefficient of determination; and τ : Kendall's tau ranking accuracy metric. Subscripts and superscripts indicate 95% confidence intervals on these statistics.

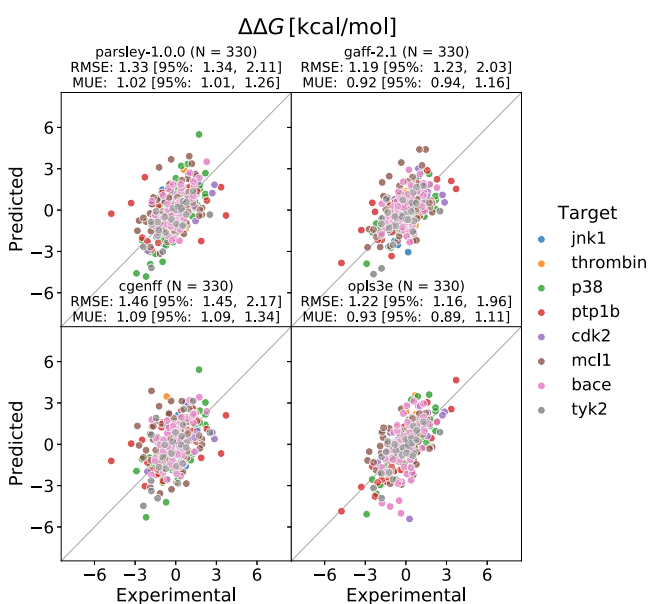


Figure 15. Results of protein–ligand binding free energy benchmarks. The predicted relative binding free energies $\Delta\Delta G$ vs the experimental results for 330 alchemical perturbation calculations with four force fields: Parsley-1.0.0, GAFFv2.1, CGenFFv3.0.1, and OPLS3e. The graphs for the latter three force fields use data reported in ref 82. The different colors denote the eight different protein targets.

reference and “Parsley” thereafter. Newer Open Force Fields are in development, and updates in the OpenFF 1.x series will also be referred to as Parsley, while new major versions will receive updated codenames. To cite Parsley, please reference the latest version of this article and the DOI of the force field version you use. This information is available in the OpenForceField repository,⁹³ and the present version may be cited as ref 44.

To provide feedback on the performance of the OpenFF force fields, we highly recommend using the issue tracker at <http://github.com/openforcefield/openforcefields>. For toolkit feed-

back, use <http://github.com/openforcefield/openforcefield>. Alternatively, inquiries may be e-mailed to support@openforcefield.org, although responses to e-mails sent to this address may be delayed, and GitHub issues receive higher priority. For information on getting started with OpenFF, please see the documentation linked at <http://github.com/openforcefield/openforcefield>, and note the availability of several introductory examples.

5. CONCLUSIONS AND DIRECTIONS

We have described a methodology to derive new simulation force fields and an initial application of this infrastructure to create OpenFF 1.0.0 (codenamed Parsley), a SMIRNOFF force field with bonded terms optimized against a range of gas-phase QM reference data. For both training and test sets, Parsley provides more accurate molecular geometries and conformational energetics, while preserving accuracy for a range of condensed phase properties. Importantly, it also yields highly competitive accuracy in calculations of relative protein–ligand binding free energies. This work lays a foundation for efficient iterative force field improvement, which is already underway in subsequent releases (OpenFF 1.1, 1.2, 1.2.1, and 1.3,^{37–41} to be described elsewhere). This work could also be naturally extended to automatically derive molecule-specific (i.e., “bespoke”) parameters,⁹⁴ and we are actively developing software tools for this application.

In the near term, we aim to extend the optimization to nonbonded interaction parameters and incorporate expanded training and testing data sets. Later, we plan to address issues related to the definitions of chemical types and elaboration of the functional form, such as by the addition of off-center partial charges and incorporation of an explicit treatment of electronic polarizability. At the same time, we hope that the associated open release of our data sets and infrastructure will enable independent use of these data and tools to advance force field science.

■ ASSOCIATED CONTENT**SI Supporting Information**

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.1c00571>.

Links to scripts and resources for training, testing, and running simulations using the Parsley force field; simulation procedures for physical properties of organic liquids; and optimized geometry RMSEs for the full test set grouped by the parameter type (PDF)

■ AUTHOR INFORMATION**Corresponding Authors**

Michael R. Shirts – Chemical & Biological Engineering Department, The University of Colorado at Boulder, Boulder, Colorado 80309, United States; orcid.org/0000-0003-3249-1097; Email: michael.shirts@colorado.edu

Michael K. Gilson – Skaggs School of Pharmacy and Pharmaceutical Sciences, The University of California at San Diego, La Jolla, California 92093, United States; orcid.org/0000-0002-3375-1738; Email: mgilson@health.ucsd.edu

John D. Chodera – Computational & Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, New York 10065, United States; orcid.org/0000-0003-0542-119X; Email: john.chodera@choderalab.org

David L. Mobley – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States; orcid.org/0000-0002-1083-5533; Email: dmobley@uci.edu

Lee-Ping Wang – Chemistry Department, The University of California at Davis, Davis, California 95616, United States; orcid.org/0000-0003-3072-9946; Email: leeping@ucdavis.edu

Authors

Yudong Qiu – Chemistry Department, The University of California at Davis, Davis, California 95616, United States; orcid.org/0000-0003-4345-8356

Daniel G. A. Smith – The Molecular Sciences Software Institute (MolSSI), Blacksburg, Virginia 24060, United States; orcid.org/0000-0001-8626-0900

Simon Boothroyd – Computational & Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, New York 10065, United States

Hyesu Jang – Chemistry Department, The University of California at Davis, Davis, California 95616, United States

David F. Hahn – Computational Chemistry, Janssen Research & Development, Beerse B-2340, Belgium; orcid.org/0000-0003-2830-6880

Jeffrey Wagner – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States

Caitlin C. Bannan – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States; Skaggs School of Pharmacy and Pharmaceutical Sciences, The University of California at San Diego, La Jolla, California 92093, United States; orcid.org/0000-0003-2777-1174

Trevor Gokey – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States

Victoria T. Lim – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States; orcid.org/0000-0003-4030-9312

Chaya D. Stern – Computational & Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, New York 10065, United States

Andrea Rizzi – Computational & Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, New York 10065, United States; Tri-Institutional Training Program in Computational Biology and Medicine, New York, New York 10065, United States

Bryon Tjanaka – Chemistry Department, The University of California at Irvine, Irvine, California 92617, United States

Gary Tresadern – Computational Chemistry, Janssen Research & Development, Beerse B-2340, Belgium; orcid.org/0000-0002-4801-1644

Xavier Lucas – F. Hoffmann-La Roche AG, Basel 4070, Switzerland; orcid.org/0000-0003-0887-3976

Christopher I. Bayly – OpenEye Scientific Software, Santa Fe, New Mexico 87508, United States

Complete contact information is available at:

<https://pubs.acs.org/doi/10.1021/acs.jctc.1c00571>

Notes

The authors declare the following competing financial interest(s): The Chodera laboratory receives or has received funding from multiple sources, including the National Institutes of Health, the National Science Foundation, the Parker Institute for Cancer Immunotherapy, Relay Therapeutics, Entasis Therapeutics, Silicon Therapeutics, EMD Serono (Merck KGaA), AstraZeneca, Vir Biotechnology, XtalPi, Foresite Labs, the Molecular Sciences Software Institute, the Starr Cancer Consortium, the Open Force Field Consortium, Cycle for Survival, a Louis V. Gerstner Young Investigator Award, and the Sloan Kettering Institute. A complete funding history for the Chodera lab can be found at <http://choderalab.org/funding>. JDC is a current member of the Scientific Advisory Board of OpenEye Scientific Software and scientific consultant to Foresite Labs. MRS is an Open Science Fellow for Silicon Therapeutics. MKG has an equity interest in and is a cofounder and scientific advisor of VeraChem LLC. DLM is a current member of the Scientific Advisory Board of OpenEye Scientific Software and an Open Science Fellow for Silicon Therapeutics.

■ ACKNOWLEDGMENTS

In this work, we stand on the shoulders of giants in the force field development community, and it would be impossible to provide a complete list of all those whose work we have benefitted from. However, we are particularly indebted to the AMBER force field community, from which we derived the initial small-molecule force fields, which provided the starting point for this work. We thank the Open Force Field Consortium for funding, including our industry partners as listed at the Open Force Field website.⁹⁵ We gratefully acknowledge Owen Madin and Karmen Condic-Jurkic for helpful discussions, along with all current and former members of the Open Force Field Initiative and the Open Force Field Scientific Advisory Board. We thank Vytautas Gapsys for assistance with the binding free energy calculations and for the insightful discussions. We thank the National Institutes of Health (NIGMS R01GM132386) for funding longer term aspects of this initiative. M.K.G. acknowledges funding from the National Institute of General Medical Sciences (GM061300). L.-P.W. and Y.Q. acknowledge support from ACS PRF 58158-DNI6. M.R.S. and J.D.C. acknowledge support from NSF CHE-1738975. D.G.A.S. acknowledges funding from the U. S.

National Science Foundation (NSF) grant ACI-1547580. These findings are solely of the authors and do not necessarily represent the views of the NIH or NSF.

REFERENCES

- (1) Coleman, R. G.; Carchia, M.; Sterling, T.; Irwin, J. J.; Shoichet, B. K. Ligand Pose and Orientational Sampling in Molecular Docking. *PLoS One* **2013**, *8*, No. e75992.
- (2) Allen, W. J.; Balias, T. E.; Mukherjee, S.; Brozell, S. R.; Moustakas, D. T.; Lang, P. T.; Case, D. A.; Kuntz, I. D.; Rizzo, R. C. DOCK 6: Impact of new features and current docking performance. *J. Comput. Chem.* **2015**, *36*, 1132–1156.
- (3) Österberg, F.; Morris, G. M.; Sanner, M. F.; Olson, A. J.; Goodsell, D. S. Automated docking to multiple target structures: Incorporation of protein mobility and structural water heterogeneity in AutoDock. *Proteins: Struct., Funct., Bioinf.* **2002**, *46*, 34–40.
- (4) Friesner, R. A.; Banks, J. L.; Murphy, R. B.; Halgren, T. A.; Klicic, J. J.; Mainz, D. T.; Repasky, M. P.; Knoll, E. H.; Shelley, M.; Perry, J. K.; Shaw, D. E.; Francis, P.; Shenkin, P. S. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy. *J. Med. Chem.* **2004**, *47*, 1739–1749 pMID: 15027865.
- (5) Jain, A. N. Surflex: Fully Automatic Flexible Molecular Docking Using a Molecular Similarity-Based Search Engine. *J. Med. Chem.* **2003**, *46*, 499–511 pMID: 12570372.
- (6) Jones, G.; Willett, P.; Glen, R. C.; Leach, A. R.; Taylor, R. Development and validation of a genetic algorithm for flexible docking. *J. Mol. Biol.* **1997**, *267*, 727–748.
- (7) McGann, M. R.; Almond, H. R.; Nicholls, A.; Grant, J. A.; Brown, F. K. Gaussian docking functions. *Biopolymers* **2003**, *68*, 76–90.
- (8) Lamb, M. L.; Jorgensen, W. L. Computational approaches to molecular recognition. *Curr. Opin. Chem. Biol.* **1997**, *1*, 449–457.
- (9) Gilson, M. K.; Zhou, H.-X. Calculation of Protein-Ligand Binding Affinities. *Annu. Rev. Biophys. Biomol. Struct.* **2007**, *36*, 21–42 pMID: 17201676.
- (10) Gallicchio, E.; Levy, R. M. Recent theoretical and computational advances for modeling protein–ligand binding affinities. In *Computational Chemistry Methods in Structural Biology*; Christov, C., Ed.; Advances in Protein Chemistry and Structural Biology; Academic Press, 2011; Vol. 85, pp 27–80.
- (11) 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. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2019**, *9*, No. e1393.
- (12) Wang, L.; Wu, Y.; Deng, Y.; Kim, B.; Pierce, L.; Krilov, G.; Lupyan, D.; Robinson, S.; Dahlgren, M. K.; Greenwood, J.; Romero, D. L.; Masse, C.; Knight, J. L.; Steinbrecher, T.; Beuming, T.; Damm, W.; Harder, E.; Sherman, W.; Brewer, M.; Wester, R.; et al. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *J. Am. Chem. Soc.* **2015**, *137*, 2695–2703 pMID: 25625324.
- (13) Cournia, Z.; Allen, B.; Sherman, W. Relative Binding Free Energy Calculations in Drug Discovery: Recent Advances and Practical Considerations. *J. Chem. Inf. Model.* **2017**, *57*, 2911–2937 pMID: 29243483.
- (14) Schindler, C.; Baumann, H.; Blum, A.; Böse, D.; Buchstaller, H. P.; Burgdorf, L.; Cappel, D.; Chekler, E.; Czodrowski, P.; Dorsch, D.; Eguida, M.; Follows, B.; Fuch β , T.; Grädler, U.; Gunera, J.; Johnson, T.; Lebrun, C. J.; Karra, S.; Klein, M.; Kötzner, L.; et al. *Large-Scale Assessment of Binding Free Energy Calculations in Active Drug Discovery Projects*, 2020. https://chemrxiv.org/articles/preprint/Large-Scale_Assessment_of_Binding_Free_Energy_Calculations_in_Active_Drug_Discovery_Projects/11364884. 10.26434/chemrxiv.11364884.v2.
- (15) Cournia, Z.; Allen, B. K.; Beuming, T.; Pearlman, D. A.; Radak, B. K.; Sherman, W. Rigorous Free Energy Simulations in Virtual Screening. *J. Chem. Inf. Model.* **2020**, *60*, 4153–4169 pMID: 32539386.
- (16) Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Gould, I. R.; Merz, K. M.; Ferguson, D. M.; Spellmeyer, D. C.; Fox, T.; Caldwell, J. W.; Kollman, P. A. A Second Generation Force Field for the Simulation of Proteins, Nucleic Acids, and Organic Molecules. *J. Am. Chem. Soc.* **1995**, *117*, 5179–5197.
- (17) MacKerell, A. D.; Bashford, D.; Bellott, M.; Dunbrack, R. L.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F. T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; et al. All-Atom Empirical Potential for Molecular Modeling and Dynamics Studies of Proteins. *J. Phys. Chem. B* **1998**, *102*, 3586–3616 pMID: 24889800.
- (18) Foloppe, N.; MacKerell, A. D., Jr. All-atom empirical force field for nucleic acids: I. Parameter optimization based on small molecule and condensed phase macromolecular target data. *J. Comput. Chem.* **2000**, *21*, 86–104.
- (19) MacKerell, A. D., Jr.; Banavali, N. K. All-atom empirical force field for nucleic acids: II. Application to molecular dynamics simulations of DNA and RNA in solution. *J. Comput. Chem.* **2000**, *21*, 105–120.
- (20) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins: Struct., Funct., Bioinf.* **2006**, *65*, 712–725.
- (21) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Struct., Funct., Bioinf.* **2010**, *78*, 1950–1958.
- (22) Wang, L.-P.; McKiernan, K. A.; Gomes, J.; Beauchamp, K. A.; Head-Gordon, T.; Rice, J. E.; Swope, W. C.; Martinez, T. J.; Pande, V. S. Building a More Predictive Protein Force Field: A Systematic and Reproducible Route to AMBER-FB15. *J. Phys. Chem. B* **2017**, *121*, 4023–4039 pMID: 28306259.
- (23) Oostenbrink, C.; Villa, A.; Mark, A. E.; Van Gunsteren, W. F. A biomolecular force field based on the free enthalpy of hydration and solvation: The GROMOS force-field parameter sets 53A5 and 53A6. *J. Comput. Chem.* **2004**, *25*, 1656–1676.
- (24) Shi, Y.; Xia, Z.; Zhang, J.; Best, R.; Wu, C.; Ponder, J. W.; Ren, P. Polarizable Atomic Multipole-Based AMOEBA Force Field for Proteins. *J. Chem. Theory Comput.* **2013**, *9*, 4046–4063 pMID: 24163642.
- (25) Kim, S.; Thiessen, P. A.; Bolton, E. E.; Chen, J.; Fu, G.; Gindulyte, A.; Han, L.; He, J.; He, S.; Shoemaker, B. A.; Wang, J.; Yu, B.; Zhang, J.; Bryant, S. H. PubChem Substance and Compound databases. *Nucleic Acids Res.* **2016**, *44*, D1202–D1213.
- (26) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and testing of a general amber force field. *J. Comput. Chem.* **2004**, *25*, 1157–1174.
- (27) Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, L.; Mackerell, A. D. CHARMM general force field: A force field for drug-like molecules compatible with the CHARMM all-atom additive biological force fields. *J. Comput. Chem.* **2010**, *31*, 671.
- (28) 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. *J. Am. Chem. Soc.* **1996**, *118*, 11225–11236.
- (29) 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. *J. Phys. Chem. B* **2014**, *118*, 6438–6446 pMID: 24702668.
- (30) Henriksen, N. M.; Fenley, A. T.; Gilson, M. K. Computational Calorimetry: High-Precision Calculation of Host–Guest Binding Thermodynamics. *J. Chem. Theory Comput.* **2015**, *11*, 4377–4394 pMID: 26523125.

- (31) Yin, J.; Henriksen, N. M.; Muddana, H. S.; Gilson, M. K. Bind3P: Optimization of a Water Model Based on Host–Guest Binding Data. *J. Chem. Theory Comput.* **2018**, *14*, 3621–3632 pMID: 29874074.
- (32) Kenney, I. M.; Beckstein, O.; Iorga, B. I. Prediction of cyclohexane-water distribution coefficients for the SAMPL5 data set using molecular dynamics simulations with the OPLS-AA force field. *J. Comput.-Aided Mol. Des.* **2016**, *30*, 1045–1058.
- (33) Rizzi, A.; Murkli, S.; McNeill, J. N.; Yao, W.; Sullivan, M.; Gilson, M. K.; Chiu, M. W.; Isaacs, L.; Gibb, B. C.; Mobley, D. L.; Chodera, J. D. Overview of the SAMPL6 host–guest binding affinity prediction challenge. *J. Comput.-Aided Mol. Des.* **2018**, *32*, 937–963.
- (34) Roos, K.; Wu, C.; Damm, W.; Reboul, M.; Stevenson, J. M.; Lu, C.; Dahlgren, M. K.; Mondal, S.; Chen, W.; Wang, L.; Abel, R.; Friesner, R. A.; Harder, E. D. OPLS3e: Extending Force Field Coverage for Drug-Like Small Molecules. *J. Chem. Theory Comput.* **2019**, *15*, 1863–1874.
- (35) Mobley, D. L.; Bannan, C. C.; Rizzi, A.; Bayly, C. I.; Chodera, J. D.; Lim, V. T.; Lim, N. M.; Beauchamp, K. A.; Slochow, D. R.; Shirts, M. R.; Gilson, M. K.; Eastman, P. K. Escaping Atom Types in Force Fields Using Direct Chemical Perception. *J. Chem. Theory Comput.* **2018**, *14*, 6076–6092 pMID: 30351006.
- (36) Bayly, C.; McKay, D.; Truchon, J. An Informal AMBER Small Molecule Force Field; Parm@Frosst, 2010. http://www.ccl.net/ccaf/data/parm_at_Frosst/.
- (37) Jang, H. Update on Parsley Minor Releases (Openff-1.1.0, 1.2.0), 2020. 10.5281/zenodo.3781313.
- (38) Jang, H.; Maat, J.; Qiu, Y.; Smith, D. G. A.; Boothroyd, S.; Wagner, J.; Bannan, C. C.; Gokey, T.; Lim, V. T.; Lucas, X.; Tjanaka, B.; Shirts, M. R.; Gilson, M. K.; Chodera, J. D.; Bayly, C. I.; Mobley, D. L.; Wang, L. P. Openforcefield/Openforcefields: Version 1.2.0; "Parsley" Update; Zenodo, 2020. 10.5281/zenodo.3872244.
- (39) Jang, H.; Maat, J.; Qiu, Y.; Smith, D. G. A.; Boothroyd, S.; Wagner, J.; Bannan, C. C.; Gokey, T.; Lim, V. T.; Lucas, X.; Tjanaka, B.; Shirts, M. R.; Gilson, M. K.; Chodera, J. D.; Bayly, C. I.; Mobley, D. L.; Wang, L. P. Openforcefield/Openforcefields: Version 1.1.0 "Parsley" Update; Zenodo, 2020. 10.5281/zenodo.3695094.
- (40) Wagner, J.; Thompson, M.; Dotson, D.; Rodríguez-Guerra, J. Openforcefield/Openforcefields: Version 1.2.1 "Parsley" Update; Zenodo, 2020. 10.5281/zenodo.4021623.
- (41) Wagner, J.; Thompson, M.; Dotson, D.; Rodríguez-Guerra, J. Openforcefield/Openforcefields: Version 1.3.0 "Parsley" Update; Zenodo, 2020. 10.5281/zenodo.4118484.
- (42) Wang, J.; Cieplak, P.; Kollman, P. A. How well does a restrained electrostatic potential (RESP) model perform in calculating conformational energies of organic and biological molecules? *J. Comput. Chem.* **2000**, *21*, 1049–1074.
- (43) Wang, L.-P.; Martinez, T. J.; Pande, V. S. Building Force Fields: An Automatic, Systematic, and Reproducible Approach. *J. Phys. Chem. Lett.* **2014**, *5*, 1885–1891 pMID: 26273869.
- (44) Wagner, J. Openforcefield/Openforcefields: Version 1.0.0 "Parsley"; Zenodo, 2019. 10.5281/zenodo.3483227.
- (45) Vilar, S.; Cozza, G.; Moro, S. Medicinal Chemistry and the Molecular Operating Environment (MOE): Application of QSAR and Molecular Docking to Drug Discovery. *Curr. Top. Med. Chem.* **2008**, *8*, 1555–1572.
- (46) eMolecules Plus Database Download, 2013. <https://reaxys.emolecules.com/index.php> (Accessed June, 2019).
- (47) MolSSI QCFractal Documentation, 2019. <https://qcfractal.readthedocs.io> (Accessed June, 2019).
- (48) MolSSI QCArchive Web Page, 2020. <https://qcarchive.molssi.org/> (Accessed June, 2019).
- (49) Smith, D. G. A.; Altarawy, D.; Burns, L. A.; Welborn, M.; Naden, L. N.; Ward, L.; Ellis, S.; Pritchard, B. P.; Crawford, T. D. The MolSSI QCArchive project: An open-source platform to compute, organize, and share quantum chemistry data. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2021**, *11*, No. e1491.
- (50) Godbout, N.; Salahub, D. R.; Andzelm, J.; Wimmer, E. Optimization of Gaussian-type basis sets for local spin density functional calculations. Part I. Boron through neon, optimization technique and validation. *Can. J. Chem.* **1992**, *70*, 560–571.
- (51) Becke, A. D. Density-functional thermochemistry. III. The role of exact exchange. *J. Chem. Phys.* **1993**, *98*, 5648–5652.
- (52) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu. *J. Chem. Phys.* **2010**, *132*, 154104.
- (53) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *J. Comput. Chem.* **2011**, *32*, 1456–1465.
- (54) Rezáč, J.; Bím, D.; Gutten, O.; Rulišek, L. Toward Accurate Conformational Energies of Smaller Peptides and Medium-Sized Macrocycles: MPCONF196 Benchmark Energy Data Set. *J. Chem. Theory Comput.* **2018**, *14*, 1254–1266 pMID: 29461829.
- (55) Kesharwani, M. K.; Karton, A.; Martin, J. M. L. Benchmark ab Initio Conformational Energies for the Proteinogenic Amino Acids through Explicitly Correlated Methods. Assessment of Density Functional Methods. *J. Chem. Theory Comput.* **2016**, *12*, 444–454 pMID: 26653705.
- (56) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.* **2005**, *7*, 3297–3305.
- (57) Tian, C.; Kasavajhala, K.; Belfon, K. A. A.; Raguette, L.; Huang, H.; Miguess, A. N.; Bickel, J.; Wang, Y.; Pincay, J.; Wu, Q.; Simmerling, C. ff19SB: Amino-Acid-Specific Protein Backbone Parameters Trained against Quantum Mechanics Energy Surfaces in Solution. *J. Chem. Theory Comput.* **2020**, *16*, 528–552.
- (58) Stern, C. D.; Bayly, C. I.; Smith, D. G. A.; Fass, J.; Wang, L.-P.; Mobley, D. L.; Chodera, J. D. Capturing non-local through-bond effects when fragmenting molecules for quantum chemical torsion scans, 2020. bioRxiv:10.1101/2020.08.27.270934.
- (59) Wang, L.-P.; Song, C. Geometry optimization made simple with translation and rotation coordinates. *J. Chem. Phys.* **2016**, *144*, 214108.
- (60) Parrish, R. M.; Burns, L. A.; Smith, D. G. A.; Simmonett, A. C.; DePrince, A. E.; Hohenstein, E. G.; Bozkaya, U.; Sokolov, A. Y.; Di Remigio, R.; Richard, R. M.; Gonthier, J. F.; James, A. M.; McAlexander, H. R.; Kumar, A.; Saitow, M.; Wang, X.; Pritchard, B. P.; Verma, P.; Schaefer, H. F.; Patkowski, K.; et al. Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability. *J. Chem. Theory Comput.* **2017**, *13*, 3185–3197 pMID: 28489372.
- (61) Smith, D. G. A.; Burns, L. A.; Simmonett, A. C.; Parrish, R. M.; Schieber, M. C.; Galvelis, R.; Kraus, P.; Kruse, H.; Di Remigio, R.; Alenaizan, A.; James, A. M.; Lehtola, S.; Misiewicz, J. P.; Scheurer, M.; Shaw, R. A.; Schriber, J. B.; Xie, Y.; Glick, Z. L.; Sirianni, D. A.; O'Brien, J. S.; et al. Psi4 1.4: Open-source software for high-throughput quantum chemistry. *J. Chem. Phys.* **2020**, *152*, 184108.
- (62) Mobley, D. L.; Wagner, J.; Chodera, J.; Bannan, C.; Rizzi, A.; Camila, Bayly, C.; SimonBoothroyd; Lim, N. M.; Lim, V.; Zhao, Y.; Lee-Ping Openforcefield/Openforcefield: 0.4.1 Bugfix Release; Zenodo, 2019. 10.5281/zenodo.3266364.
- (63) Qiu, Y.; Smith, D. G. A.; Stern, C. D.; Feng, M.; Jang, H.; Wang, L.-P. Driving torsion scans with wavefront propagation. *J. Chem. Phys.* **2020**, *152*, 244116.
- (64) Paton, R. S.; Goodman, J. M. Hydrogen Bonding and Pi-Stacking: How Reliable are Force Fields? A Critical Evaluation of Force Field Descriptions of Nonbonded Interactions. *J. Chem. Inf. Model.* **2009**, *49*, 944–955.
- (65) McGibbon, R. T.; Beauchamp, K. A.; Harrigan, M. P.; Klein, C.; Swails, J. M.; Hernández, C. X.; Schwantes, C. R.; Wang, L.-P.; Lane, T. J.; Pande, V. S. MDTraj: A Modern Open Library for the Analysis of Molecular Dynamics Trajectories. *Biophys. J.* **2015**, *109*, 1528–1532.
- (66) Harder, E.; Damm, W.; Maple, J.; Wu, C.; Reboul, M.; Xiang, J. Y.; Wang, L.; Lupyan, D.; Dahlgren, M. K.; Knight, J. L.; et al. OPLS3: a force field providing broad coverage of drug-like small molecules and proteins. *J. Chem. Theory Comput.* **2016**, *12*, 281–296.

- (67) Cerutti, D. S.; Swope, W. C.; Rice, J. E.; Case, D. A. ff14ipq: A self-consistent force field for condensed-phase simulations of proteins. *J. Chem. Theory Comput.* **2014**, *10*, 4515–4534.
- (68) Zgarbová, M.; Otyepka, M.; Šponer, J.; Mládek, A.; Banáš, P.; Cheatham, T. E., III; Jurečka, P. Refinement of the Cornell et al. nucleic acids force field based on reference quantum chemical calculations of glycosidic torsion profiles. *J. Chem. Theory Comput.* **2011**, *7*, 2886–2902.
- (69) 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.
- (70) Development Version of ForceBalance Based on v1.6.0 Used for 1.0.0 Valence Parameter Fitting, 2019. <https://github.com/leeping/forcebalance/tree/Sb3a65d1baefe17625c9b05c0db5d27458c5344a>.
- (71) Lim, V. T.; Hahn, D. F.; Tresadern, G.; Bayly, C. I.; Mobley, D. Benchmark Assessment of Molecular Geometries and Energies from Small Molecule Force Fields, 2020. chemRxiv:10.26434/chemrxiv.12551867.v2.
- (72) Ehrman, J. N.; Bannan, C. C.; Lim, V. T.; Thi, N.; Kyu, D. Y.; Mobley, D. L. Improving Force Fields by Identifying and Characterizing Small Molecules with Parameter Inconsistencies; Zenodo, 2019.
- (73) Ehrman, J. N.; Lim, V. T.; Bannan, C. C.; Thi, N.; Kyu, D. Y.; Mobley, D. L. Improving Small Molecule Force Fields by Identifying and Characterizing Small Molecules with Inconsistent Parameters. *J. Comput.-Aided Mol. Des.* **2021**, *35*, 271–284.
- (74) Sure, R.; Grimme, S. Corrected small basis set Hartree-Fock method for large systems. *J. Comput. Chem.* **2013**, *34*, 1672–1685.
- (75) Sresht, V.; Rai, B. Calculation of Torsional Strain Energy; PfizerRD Torsional-Strain, 2019. <https://github.com/PfizerRD/torsional-strain>.
- (76) Sterling, T.; Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *J. Chem. Inf. Model.* **2015**, *55*, 2324–2337 pMID: 26479676.
- (77) QCArchive Basic Examples, Optimization Data sets. https://qcarcivetutorials.readthedocs.io/en/latest/basic_examples/optimization_datasets.html.
- (78) Frenkel, M.; Chiroco, R. D.; Diky, V.; Dong, Q.; Marsh, K. N.; Dymond, J. H.; Wakeham, W. A.; Stein, S. E.; Königsberger, E.; Goodwin, A. R. H. XML-based IUPAC standard for experimental, predicted, and critically evaluated thermodynamic property data storage and capture (ThermoML) (IUPAC Recommendations 2006). *Pure Appl. Chem.* **2006**, *78*, 541–612.
- (79) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.* **1983**, *79*, 926–935.
- (80) Beauchamp, K. A.; Behr, J. M.; Rustenburg, A. S.; Bayly, C. I.; Kroenlein, K.; Chodera, J. D. Toward Automated Benchmarking of Atomistic Force Fields: Neat Liquid Densities and Static Dielectric Constants from the ThermoML Data Archive. *J. Phys. Chem. B* **2015**, *119*, 12912–12920.
- (81) Hahn, D. F.; Bayly, C. I.; Macdonald, H. E. B.; Chodera, J. D.; Mey, A. S. J. S.; Mobley, D. L.; Benito, L. P.; Schindler, C. E. M.; Tresadern, G.; Warren, G. L. Best practices for constructing, preparing, and evaluating protein-ligand binding affinity benchmarks, 2021. arXiv: 2105.06222.
- (82) 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. *Chem. Sci.* **2020**, *11*, 1140–1152.
- (83) Hahn, D. F.; Wagner, J. Openforcefield/Protein–Ligand-Benchmark: 0.1.2. Release to Create Zenodo Record; Zenodo, 2021. 10.5281/ZENODO.4813735.
- (84) Seeliger, D.; de Groot, B. L. Protein Thermostability Calculations Using Alchemical Free Energy Simulations. *Biophys. J.* **2010**, *98*, 2309–2316.
- (85) Gapsys, V.; Michielssens, S.; Seeliger, D.; de Groot, B. L. pmx: Automated protein structure and topology generation for alchemical perturbations. *J. Comput. Chem.* **2015**, *36*, 348–354.
- (86) Best, R. B.; Hummer, G. Optimized Molecular Dynamics Force Fields Applied to the Helix-Coil Transition of Polypeptides. *J. Phys. Chem. B* **2009**, *113*, 9004–9015.
- (87) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field: Improved Protein Side-Chain Potentials. *Proteins: Struct., Funct., Bioinf.* **2010**, *78*, 1950–1958.
- (88) Hahn, D. F. Protein–Ligand-Benchmark-Analysis: 0.1.0 Initial Release; Zenodo, 2021. 10.5281/ZENODO.4894294.
- (89) Bruce Macdonald, H. E. Openforcefield/Arsenic. Open Force Field Initiative, 2020.
- (90) Hahn, D. F.; Tresadern, G. J.; Mobley, D. L. Large Scale Benchmarking the Prospective Accuracy of Protein–Ligand Free Energy Calculations, version number: 1.0.0.; Zenodo, 2020. 10.5281/ZENODO.4674371.
- (91) Qiu, Y. OpenFF “parsley” 1.0.0-RC2 Release Package, 2019. <https://github.com/openforcefield/openforcefield-forcebalance/releases/tag/v1.0.0-RC2>.
- (92) Jang, H.; Qiu, Y. Openforcefield Parsley 1.0.0 Benchmark Release Package, 2019. <https://github.com/openforcefield/release-1-benchmarking/releases/tag/v1.0.0>.
- (93) Wagner, J.; Thompson, M.; Dotson, D.; Jang, H.; Rodríguez-Guerra, J. Repository Containing Force Fields Released by the Open Force Field Initiative, 2020. <https://github.com/openforcefield/openforcefields>.
- (94) Horton, J. T.; Allen, A. E. A.; 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.
- (95) OpenForceField, about/organization/, 2020. <https://openforcefield.org/about/organization/> (Accessed October, 2020).