



HHS Public Access

Author manuscript

J Phys Chem B. Author manuscript; available in PMC 2016 October 08.

Published in final edited form as:

J Phys Chem B. 2015 October 8; 119(40): 12912–12920. doi:10.1021/acs.jpcc.5b06703.

Towards Automated Benchmarking of Atomistic Forcefields: Neat Liquid Densities and Static Dielectric Constants from the ThermoML Data Archive

Kyle A. Beauchamp^{*,†}, Julie M. Behr[‡], Ariën S. Rustenburg[¶], Christopher I. Bayly[§], Kenneth Kroenlein^{||}, and John D. Chodera^{*,†}

[†]Computational Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, NY

[‡]Tri-Institutional Program in Computational Biology and Medicine, Weill Cornell Medical College, New York, NY

[¶]Graduate Program in Physiology, Biophysics, and Systems Biology, Weill Cornell Medical College, New York, NY

[§]OpenEye Scientific Software Inc., Santa Fe, NM

^{||}Thermodynamics Research Center, NIST, Boulder, CO

Abstract

Atomistic molecular simulations are a powerful way to make quantitative predictions, but the accuracy of these predictions depends entirely on the quality of the forcefield employed. While experimental measurements of fundamental physical properties offer a straightforward approach for evaluating forcefield quality, the bulk of this information has been tied up in formats that are not machine-readable. Compiling benchmark datasets of physical properties from non-machine-readable sources requires substantial human effort and is prone to the accumulation of human errors, hindering the development of reproducible benchmarks of forcefield accuracy. Here, we examine the feasibility of benchmarking atomistic forcefields against the NIST ThermoML data archive of physicochemical measurements, which aggregates thousands of experimental measurements in a portable, machine-readable, self-annotating IUPAC-standard format. As a proof of concept, we present a detailed benchmark of the generalized Amber small molecule forcefield (GAFF) using the AM1-BCC charge model against experimental measurements (specifically bulk liquid densities and static dielectric constants at ambient pressure) automatically extracted from the archive, and discuss the extent of data available for use in larger scale (or continuously performed) benchmarks. The results of even this limited initial benchmark highlight

^{*}To whom correspondence should be addressed: kyle.beauchamp@choderalab.org; john.chodera@choderalab.org.

Supporting Information

Five additional figures and one command line instruction are provided in the supporting information. This material is available free of charge via the Internet at <http://pubs.acs.org/>.

Disclaimers

This contribution of the National Institute of Standards and Technology (NIST) is not subject to copyright in the United States. Products or companies named here are cited only in the interest of complete technical description, and neither constitute nor imply endorsement by NIST or by the U.S. government. Other products may be found to serve as well.

a general problem with fixed-charge forcefields in the representation low dielectric environments such as those seen in binding cavities or biological membranes.

1 Introduction

Recent advances in hardware and software for molecular dynamics simulation now permit routine access to atomistic simulations at the 100 ns timescale and beyond.¹ Leveraging these advances in combination with consumer GPU clusters, distributed computing, or custom hardware has brought microsecond and millisecond simulation timescales within reach of many laboratories. These dramatic advances in sampling, however, have revealed deficiencies in forcefields as a critical barrier to enabling truly predictive simulations of physical properties of biomolecular systems.

Protein and water forcefields have been the subject of numerous benchmarks^{2–4} and enhancements,^{5–7} with key outcomes including the ability to fold fast-folding proteins,^{8–10} improved fidelity of water thermodynamic properties,¹¹ and improved prediction of NMR observables. Although small molecule forcefields have also been the subject of benchmarks^{12–14} and improvements,¹⁵ such work has typically focused on small perturbations to specific functional groups. For example, a recent study found that modified hydroxyl nonbonded parameters led to improved prediction of static dielectric constants and hydration free energies.¹⁵ There are also outstanding questions of generalizability of these targeted perturbations; it is uncertain whether changes to the parameters for a specific chemical moiety will be compatible with seemingly unrelated improvements to other groups. Addressing these questions requires establishing community agreement upon shared benchmarks that can be easily replicated among laboratories to test proposed forcefield enhancements and expanded as the body of experimental data grows.

A key barrier to establishing reproducible and extensible forcefield accuracy benchmarks is that many experimental datasets are heterogeneous, paywalled, and unavailable in machine-readable formats (although notable counterexamples exist, e.g., the PDB,¹⁶ Free-Solv,¹⁷ and the BMRB¹⁸). While this inconvenience is relatively minor for benchmarking forcefield accuracy for a single target system (e.g., water), it becomes prohibitive for studies spanning relevant chemical spaces, such as forcefields intended to describe a large variety of druglike small organic molecules.

In addition to inconvenience, the number and kind of human-induced errors that can corrupt hand-compiled benchmarks are legion. A United States Geological Survey (USGS) case study examining the reporting and use of literature values of the aqueous solubility (S_w) and octanol-water partition coefficients (K_{ow}) for DDT and its persistent metabolite DDE provides remarkable insight into a variety of common errors.¹⁹ Secondary sources are often cited as primary sources—a phenomenon that occurred up to five levels deep in the case of DDT/DDE; citations for data are often incorrect, misattributed to unrelated publications, or omitted altogether; numerical data can be mistranscribed, transposed, or incorrectly converted among unit systems.¹⁹ In the case of DDT/DDE, these errors occur to such a degree that the authors note “strings of erroneous data compose as much as 41–73 percent of

the total data”.¹⁹ Given the number and importance of these measurements, the quality of physicochemical datasets of lesser importance may be suspect.

To ameliorate problems of data archival, the NIST Thermodynamics Research Center (TRC) has developed an IUPAC standard XML-based format—ThermoML^{20–22}—for storing physicochemical measurements, uncertainties, and metadata. Manuscripts containing new experimental measurements submitted to several journals (*J. Chem. Eng. Data*, *J. Chem. Thermodyn.*, *Fluid Phase Equilib.*, *Thermochemica Acta*, and *Int. J. Thermophys.*) are guided through a data archival process that involves sanity checks, conversion to a standard machine-readable format, and archival at the TRC (<http://trc.nist.gov/ThermoML.html>).

Here, we examine the ThermoML archive as a potential source for a reproducible, extensible accuracy benchmark of biomolecular forcefields. As a proof of concept, we concentrate on two important physical property measurements easily computable in many simulation codes—neat liquid density and static dielectric constant measurements—with the goal of developing a standard benchmark for validating these properties in fixed-charge forcefields of drug-like molecules and biopolymer residue analogues. These two properties provide sensitive tests of forcefield accuracy that are nonetheless straightforward to calculate. Using these data, we evaluate the generalized Amber small molecule forcefield (GAFF)^{23,24} with the AM1-BCC charge model^{25,26} and identify systematic biases to aid further forcefield refinement.

2 Methods

2.1 ThermoML Archive retrieval and processing

A snapshot of the ThermoML Archive was obtained from the the NIST TRC on 8 Apr. 2015. To explore the content of this archive, we created a Python (version 2.7.9) tool—ThermoPyL (<https://github.com/choderalab/ThermoPyL>)—that formats the XML content into a spreadsheet-like format accessible via the Pandas (version 0.15.2) library. This tool also contains a preliminary version of scripts for maintaining an up-to-date version of the ThermoML archive. First, we obtained the XML schema (<http://media.iupac.org/namespaces/ThermoML/ThermoML.xsd>) defining the layout of the data. This schema was converted into a Python object via PyXB 1.2.4 (<http://pyxb.sourceforge.net/>). Finally, this schema was used to extract the data into Pandas²⁷ dataframes, and successive data filters described in Section 3.1 were applied to explore the composition of the data.

2.2 Simulation

To enable automated accuracy benchmarking of physicochemical properties of neat liquids such as mass density and static dielectric constant, we developed a semi-automated pipeline for preparing simulations, running them on a standard computer cluster using a portable simulation package, and analyzing the resulting data. All code for this procedure is available at <https://github.com/choderalab/LiquidBenchmark>. Below, we describe the operation of the various stages of this pipeline and their application to the benchmark reported here.

2.2.1 Preparation—Chemical names were parsed from the ThermoML extract and converted to both CAS and SMILES strings using cirpy (<https://github.com/mcs07/CIRpy>).

Smiles strings were converted into molecular structures using the OpenEye Python Toolkit version 2015-2-3,²⁸ as wrapped in openmoltools.

Simulation boxes containing 1 000 molecules were constructed using PackMol version 14-225²⁹ wrapped in the Python automation library openmoltools (<http://github.com/choderalab/openmoltools/>). In order to ensure stable automated equilibration, PackMol box volumes were chosen to accommodate twice volume of the enclosed atoms, with atomic radii estimated as 1.06 Å and 1.53 Å for hydrogens and nonhydrogens, respectively.

For this illustrative benchmark, we utilized the generalized Amber small molecule force-field (GAFF)^{23,24} with the AM1-BCC charge model,^{25,26} which we shall refer to as the GAFF/AM1-BCC forcefield.

Canonical AM1-BCC^{25,26,30} charges were generated with the OpenEye Python Toolkit version 2015-2-3,²⁸ using the oequacpac.OEAssignPartialCharges module with the OECharges AM1BCCSym option, which utilizes a conformational expansion procedure (using oeomega.OEOmega³¹) prior to charge fitting to minimize artifacts from intramolecular contacts. The OEOmega selected conformer was then processed using antechamber (with parmchk2) and tleap in AmberTools 14³² to produce Amber-format prmtop and inpcrd files, which were then read into OpenMM to perform molecular simulations using the simtk.openmm.app module.

The simulations reported here used libraries openmoltools 0.6.4, OpenMM 6.3,³³ and MDTraj 1.3.³⁴ Exact commands to install various dependencies can be found in section S1.1.

2.2.2 Equilibration and production—All simulations were performed using OpenMM 6.3.³³ Simulation boxes were first minimized using the L-BFGS algorithm³⁵ using the LocalEnergyMinimizer default parameters and subsequently equilibrated for 10^7 steps with an equilibration timestep of 0.4 fs and a collision rate of 5 ps^{-1} . Production simulations employed a Langevin Leapfrog integrator³⁶ with collision rate 1 ps^{-1} and a 1 fs timestep, as we found that timesteps of 2 fs or greater led to a significant timestep dependence in computed equilibrium densities (Fig. S1).

Equilibration and production simulations utilized a Metropolis Monte Carlo barostat with a control pressure of 1 atm (101.325 kPa), utilizing molecular scaling and automated step size adjustment during equilibration, with volume moves attempted every 25 steps. The particle mesh Ewald (PME) method with conducting boundary conditions³⁷ was used with a long-range cutoff of 0.95 nm, and a long-range isotropic dispersion correction was employed to correct for the truncation of Lennard-Jones interactions outside the 0.95 nm cutoff. PME grid and spline parameters were automatically selected using the default settings in OpenMM 6.3 for the CUDA platform,³³ which was operated in default mixed-precision mode. Instantaneous densities were stored every 250 fs, while trajectory snapshots were stored every 5 ps.

Automatic termination criteria: Production simulations were continued until automatic analysis showed standard errors in densities were less than $2 \times 10^{-4} \text{ g/cm}^3$. Automatic

analysis of the production simulation data was run every 1 ns of simulation time, and utilized the detectEquilibration method³⁸ in the timeseries module of pymbar 2.1³⁹ to automatically discard the initial portion of the production simulation containing strong far-from-equilibrium configurations. This procedure selects the equilibration endpoint T_{eq} by maximizing the number of effectively uncorrelated samples in the remainder of the production simulation, $N_{\text{eff}} = (T - T_{\text{equil}})/g$, with T the total simulation length. The statistical inefficiency g was determined by autocorrelation analysis using the fast adaptive statistical inefficiency computation method as implemented in the timeseries.computeStatisticalInefficiency method of pymbar 2.1 (where the algorithm is described in⁴⁰). This approach is essentially the same as the fixed-width procedure described by eq. 7.12 of ref.,⁴¹ with n^* equal to 4000 and the sequential testing correction (n^{-1} term) ignored due to the large value of n .

Statistical errors in estimated average density $\langle \rho \rangle$ were computed by the Markov chain standard error (MCSE)

$$\delta_{\rho} \approx \frac{\text{std}(\rho)}{\sqrt{N_{\text{eff}}}} \quad (1)$$

where (ρ) is the sample standard deviation of the density and N_{eff} is the number of effectively uncorrelated samples.

Using this adaptive protocol, we found starting trajectory lengths of 12000 (8000, 16000) density frames (250 fs each), discarded regions of 28 (0, 460) density frames, and statistical inefficiencies of 20 (15, 28) density frames; reported numbers indicate (median, (25% quartile, 75% quartile)).

2.3 Timings

The wall time required for a given simulation depends on the number of atoms in the simulation system (3 000–29 000 atoms), the GPU used (GTX 680 or GTX Titan), and the time required for automated termination. For butyl acrylate (21 000 atoms) on a GTX Titan, the wall-clock performance is approximately 80 ns/day. Using 80 ns/day with approximately 3 ns of production simulation corresponds to 1 hour for the production segment of the simulation and 3 hours for the fixed equilibration portion of 10^7 steps.

2.3.1 Data analysis and statistical error estimation—Trajectory analysis was performed using OpenMM 6.3³³ and MDTraj 1.3.³⁴

Mass density: Mass density ρ was computed via the relation,

$$\rho = \left\langle \frac{M}{V} \right\rangle, \quad (2)$$

where M is the total mass of all particles in the system and V is the instantaneous volume of the simulation box.

Static dielectric constants: Static dielectric constants were calculated using the dipole fluctuation approach appropriate for PME with conducting (“tin-foil”) boundary conditions,^{11,42} with the total system box dipole μ computed from trajectory snapshots using MDTraj 1.3.³⁴

$$\epsilon = 1 + \beta \frac{4\pi}{3} \frac{\langle \mu \cdot \mu \rangle - \langle \mu \rangle \cdot \langle \mu \rangle}{\langle V \rangle} \quad (3)$$

where $\beta \equiv 1/k_{\text{B}}T$ is the inverse temperature.

Computation of expectations: Expectations were estimated by computing sample means over the production simulation after discarding the initial far-from-equilibrium portion to equilibration (as described in **Automatic termination** criteria above).

Statistical uncertainties: For density uncertainties, the Markov chain standard error (MCSE) was estimated using Eq. 1. For dielectric uncertainties, in order to avoid the complexities of computing and propagating correlated errors in Eq. 3, a bootstrap procedure was employed: the portion of the production simulation not discarded to equilibration was used as input to a circular block bootstrapping procedure⁴³ with block sizes automatically selected to maximize the error.⁴⁴

2.3.2 Code availability—All code to perform data analysis and create figures for this work, as well as all intermediate data (except configurational trajectories, due to their large size), is available at <https://github.com/choderalab/LiquidBenchmark>.

3 Results

3.1 Extracting neat liquid measurements from the NIST TRC ThermoML Archive

As described in Section 2.1, we retrieved a copy of the ThermoML Archive and performed a number of sequential filtering steps to produce an ThermoML extract relevant for benchmarking forcefields describing small organic molecules. As our aim is to explore neat liquid data with functional groups relevant to biopolymers and drug-like molecules, we applied the following ordered filters, starting with all data containing density or static dielectric constants:

1. The measured sample contains only a single component (e.g., no binary mixtures)
2. The molecule contains only druglike elements (defined here as H, N, C, O, S, P, F, Cl, Br)
3. The molecule has 10 non-hydrogen atoms
4. The measurement was performed in a biophysically relevant temperature range (270 T [K] 330)
5. The measurement was performed at ambient pressure (100 P [kPa] 102)
6. Only measurements in liquid phase were retained

7. The temperature and pressure were rounded to nearby values (as described below), averaging all measurements within each group of like conditions
8. Only conditions (molecule, temperature, pressure) for which *both* density and dielectric constants were available were retained

The temperature and pressure rounding step was motivated by common data reporting variations; for example, an experiment performed at the freezing temperature of water and ambient pressure might be entered as either 101.325 kPa or 100 kPa, with a temperature of either 273 K or 273.15 K. Therefore all pressures within the range [kPa] ($100 \leq P \leq 102$) were rounded to exactly 1 atm (101.325 kPa). Temperatures were rounded to one decimal place in K.

The application of these filters (Table 1) leaves 246 conditions—where a *condition* here indicates a (molecule, temperature, pressure) tuple—for which both density and dielectric data are available. The functional groups present in the resulting dataset are summarized in Table 2; see Section 2.1 for further description of the software pipeline used.

3.2 Benchmarking GAFF/AM1-BCC against the ThermoML Archive

3.2.1 Mass density—Mass densities of bulk liquids have been widely used for parameterizing and testing force-fields, particularly the Lennard-Jones parameters representing dispersive and repulsive interactions.^{46,47} We therefore used the present ThermoML extract as a benchmark of the GAFF/AM1-BCC forcefield (Fig. 1).

Overall accuracy: Overall, the densities show reasonable accuracy, with a root-mean square (RMS) relative error over all measurements of $(3.0 \pm 0.1)\%$, especially encouraging given that this forcefield was not designed with the intention of modeling bulk liquid properties of organic molecules.^{23,24} This is reasonably consistent with previous studies reporting relative error of 4% on a different benchmark set.¹²

Temperature dependence: For a given compound, the signs of the errors typically do not change at different temperatures (Fig. 1, Fig. S4). Furthermore, the magnitudes of the error also remain largely constant (vertical lines in Fig. 1 B), although several exceptions do occur. It is possible that these systematic density offsets indicate correctable biases in forcefield parameters.

Outliers: The largest density errors occur for a number of oxygen-containing compounds: 1,4-dioxane; 2,5,8-trioxanonane; 2-aminoethanol; dimethyl carbonate; formamide; and water (Fig. S4). The absolute error on these poor predictions is on the order of 0.05 g/cm^3 , which is substantially higher than the measurement error ($\approx 0.008 \text{ g/cm}^3$; see Fig. S2).

We note that our benchmark includes a GAFF/AM1-BCC model for water due to our desire to automate benchmarks against a forcefield capable of modeling a large variety of small molecular liquids. Water—an incredibly important solvent in biomolecular systems—is generally treated with a special-purpose model (such as TIP3P⁴⁶ or TIP4P-Ew¹¹) parameterized to fit a large quantity of thermophysical data. As expected, the GAFF/AM1-BCC model performs poorly in reproducing liquid densities for this very special solvent. We

conclude that it remains highly advisable that the field continue to use specialized water models when possible.

3.2.2 Static dielectric constant

Overall accuracy: As a measure of the dielectric response, the static dielectric constant of neat liquids provides a critical benchmark of the accuracy of electrostatic treatment in forcefield models. Discussing the accuracy in terms of the ability of GAFF/AM1-BCC to reproduce the static dielectric constant is not necessarily meaningful because of the way that the solvent dielectric enters into the Coulomb potential between two point charges separated by a distance r ,

$$U(r) = \frac{q_1 q_2}{\epsilon r} \propto \frac{1}{\epsilon}. \quad (4)$$

It is evident that $1/\epsilon$ is a much more meaningful quantity to compare than directly, as a 5% error in $1/\epsilon$ will cause a 5% error in the Coulomb potential between two point charges (assuming a uniform dielectric), while a 5% error in ϵ will have a much more complex ϵ -dependent effect on the Coulomb potential. We therefore compare simulations against measurements in our ThermoML extract on the $1/\epsilon$ scale in Fig. 2.

GAFF/AM1-BCC systematically underestimates the dielectric constants of nonpolar liquids: Overall, we find the dielectric constants to be qualitatively reasonable, but with clear deviations from experiment, particularly for nonpolar liquids. This is not surprising given the complete neglect of electronic polarization which will be the dominant contribution for such liquids. In particular, GAFF/AM1-BCC systematically underestimates the dielectric constants for nonpolar liquids, with the predictions of $\epsilon \approx 1.0$ being substantially smaller than the measured $\epsilon \approx 2$. Because this deviation likely stems from the lack of an explicit treatment of electronic polarization, we used a simple empirical polarization model that computes the molecular electronic polarizability α as a sum of elemental atomic polarizability contributions.⁴⁸

From the computed molecular electronic polarizability α , an additive correction to the simulation-derived static dielectric constant accounting for the missing electronic polarizability can be computed¹¹

$$\Delta\epsilon = 4\pi N \frac{\alpha}{\langle V \rangle} \quad (5)$$

A similar polarization correction was used in the development of the TIP4P-Ew water model, where it had a minor effect¹¹ because almost all the high static dielectric constant for water comes from the configurational response of its strong dipole. However, the missing polarizability is a dominant contribution to the static dielectric constant of nonpolar organic molecules; in the case of water, the empirical atomic polarizability model predicts a dielectric correction (to ϵ) of 0.52, while 0.79 was used for the TIP4P-Ew model. Averaging this dielectric constant over all liquids in the present work leads to average polarizability corrections (to ϵ) of 0.74 ± 0.08 . Taking the dataset as a whole, we find that the relative

error in uncorrected dielectric (to ϵ) is on the order of -0.34 ± 0.02 , as compared to -0.25 ± 0.02 for the corrected dielectric.

4 Discussion

4.1 Mass densities

Our simulations have indicated the presence of systematic density biases with magnitudes larger than the measurement error. Correcting these errors may be a low-hanging fruit for future forcefield refinements. As an example of the feasibility of improved accuracy in densities, a recent three-point water model was able to recapitulate water density with errors of less than 0.005 g/cm^3 over the temperature range [280 K, 320 K].⁴⁹ This improved accuracy in density prediction was obtained alongside accurate predictions of other experimental observables, including the static dielectric constant. We suspect that such accuracy might be obtainable for GAFF-like forcefields across some portion of chemical space. A key challenge for the field is to demarcate the fundamental limit of fixed-charge forcefields for predicting orthogonal classes of experimental observables. For example, is it possible to achieve a relative density error of 10^{-4} without sacrificing accuracy of other properties such as enthalpies of vaporization? In our opinion, the best way to answer such questions is to systematically build forcefields with the goal of predicting various properties to within their known experimental uncertainties, similar to what has been done for water.^{11,49}

4.2 Dielectric constants in forcefield parameterization

A key feature of the static dielectric constant for a liquid is that, for forcefield purposes, it consists of two very different contributions, distinguished by the dependence on the fixed charges of the forcefield and dynamic motion of the molecule. One component, the electronic polarizability (which can be separately quantified through the high-frequency dielectric constant), arises from the almost-instantaneous electronic polarization in response to the external electric field. Electronic polarizability contributes a small component, generally around $\epsilon = 2$, which can be dominant for non-polar liquids but is completely neglected by the non-polarizable forcefields in common use for biomolecular simulations.

The other component arises from the dynamical response of the molecule through reorientation or conformational relaxation via nuclear motion. For small polar liquids that lack significant internal degrees of freedom—such as water—reorientation of various molecular multipoles in response to the external electric field contributes the majority of the static dielectric constant. As a result, for polar liquids, we expect the parameterized atomic charges to play a major role in the static dielectric.

Recent forcefield development has seen a resurgence of papers fitting dielectric constants during forcefield parameterization.^{15,49} However, a number of authors have pointed out potential challenges in constructing self-consistent fixed-charge forcefields.^{50,51}

Interestingly, recent work by Dill and coworkers⁵⁰ observed that, for CCl_4 , reasonable choices of point charges are incapable of recapitulating the observed dielectric of $\epsilon = 2.2$, instead producing dielectric constants in the range of $1.0 < \epsilon < 1.05$. This behavior is quite

general: fixed point charge forcefields will predict $\epsilon \approx 1$ for many nonpolar or symmetric molecules, but the measured dielectric constants are instead $\epsilon \approx 2$ (Fig. 3). While this behavior is well-known and results from missing physics of polarizability, we suspect it may have several profound consequences, which we discuss below.

Suppose, for example, that one attempts to fit forcefield parameters to match the static dielectric constants of CCl_4 , CHCl_3 , CH_2Cl_2 , and CH_3Cl . In moving from the tetrahedrally-symmetric CCl_4 to the asymmetric CHCl_3 , it suddenly becomes possible to achieve the observed dielectric constant of 4.8 by an appropriate choice of point charges. However, the model for CHCl_3 uses fixed point charges to account for *both* the permanent dipole moment and the electronic polarizability, whereas the CCl_4 model contains no treatment of polarizability. We hypothesize that this inconsistency in parameterization may lead to discontinuous jumps in physical properties in related molecular series, where symmetric molecules (e.g., benzene and CCl_4) have qualitatively different properties than closely related asymmetric molecules (e.g., toluene and CHCl_3).

How important is this effect? We expect it to be important wherever we encounter the transfer of a polar molecule (such as a peptide, native ligand, or a pharmaceutical small molecule) from a polar environment (such as the cytosol, interstitial fluid, or blood) into a non-polar environment (such as a biological membrane or non-polar binding site of an enzyme or receptor). Thus we expect this to be implicated in biological processes ranging from ligand binding to absorption and distribution within the body. To understand this conceptually, consider the transfer of a polar small-molecule transfer from the non-polar interior of a lipid bilayer to the aqueous and hence very polar cytosol.

As a real-world example, we imagine that the missing atomic polarizability could be important in accurate transfer free energies involving low-dielectric solvents, such as the small-molecule transfer free energy from octanol or cyclohexane to water. The Onsager model for solvation of a dipole μ of radius a gives us a way to estimate the magnitude of error introduced by making an error of $\Delta\epsilon$ in the static dielectric constant of a solvent. The free energy of dipole solvation is given by this model as

$$\Delta G = -\frac{\mu^2 \epsilon - 1}{a^3 2\epsilon + 1} \quad (6)$$

such that, for an error of $\Delta\epsilon$ departing from the true static dielectric constant ϵ , we find the error in solvation is

$$\Delta\Delta G = -\frac{\mu^2}{a^3} \left[\frac{(\epsilon + \Delta\epsilon) - 1}{2(\epsilon + \Delta\epsilon) + 1} - \frac{\epsilon - 1}{2\epsilon + 1} \right] \quad (7)$$

For example, the solvation of water ($a = 1.93 \text{ \AA}$, $\mu = 2.2 \text{ D}$) in a low dielectric medium such as tetrachloromethane or benzene ($\epsilon \sim 2.2$, but $\epsilon = -1.2$) gives an error of $\Delta G \sim -8 \text{ kJ/mol}$ (-2 kcal/mol).

Implications for transfer free energies—As another example, consider the transfer of small druglike molecules from a nonpolar solvent (such as cyclohexane) to water, a property

often measured to indicate the expected degree of lipophilicity of a compound. To estimate the magnitude of error expected, for each molecule in the FreeSolv database (as of 20 Feb 2015),^{17,54} we estimated the expected error in computed transfer free energies should GAFF/AM1-BCC be used to model the nonpolar solvent cyclohexane using the Onsager model (Eq. 7). We took the cavity radius a to be half the maximum interatomic distance and calculated $\mu = \sum_i q_i r_i$ using the provided mol2 coordinates and AM1-BCC charges. This calculation predicts a mean error of (-3.8 ± 0.3) kJ/mol $[(-0.91 \pm 0.07)$ kcal/mol] for the 643 molecules (where the standard error is computed from bootstrapping over FreeSolv compound measurements), suggesting that the missing atomic polarizability unrepresentable by fixed point charge forcefields could contribute substantially to errors in predicted transfer and solvation properties of druglike molecules. In other words, the use of a fixed-charge physics may lead to errors of 3.8 kJ/mol in cyclohexane transfer free energies. We conjecture that this missing physics will be important in the upcoming (2015) SAMPL challenge,⁵⁵ which will examine transfer free energies in several low dielectric media.

Utility in parameterization—Given their ease of measurement and direct connection to long-range electrostatic interactions, static dielectric constants have high potential utility as primary data for forcefield parameterization efforts. Although this will require the use of forcefields with explicit treatment of atomic polarizability, the inconsistency of fixed-charge models in low-dielectric media is sufficiently alarming to motivate further study of polarizable forcefields. In particular, continuum methods,^{56–58} point dipole methods,^{59,60} and Drude methods^{61,62} have been maturing rapidly. Finding the optimal balance of accuracy and performance remains an open question; however, the use of experimentally-parameterized direct polarization methods⁶³ may provide polarizability physics at a cost not much greater than fixed charge forcefields. This is not to say that we suggest an immediate transition to polarizable force fields—the efficiency benefits and pervasiveness of fixed-charge models are important. Furthermore, empirical corrections such as over-polarized charges⁵⁰ and others⁵¹ may provide the ability to model some low-dielectric behaviors at the cost of some transferability.

4.3 ThermoML as a data source

The present work has focused on the neat liquid density and dielectric measurements present in the ThermoML Archive^{20,21,64} as a target for molecular dynamics forcefield validation. While liquid mass densities and static dielectric constants have already been widely used in forcefield work, several aspects of ThermoML make it a unique resource for the force-field community. First, the aggregation, support, and dissemination of ThermoML datasets through the ThermoML Archive is supported by NIST, whose mission makes these tasks a long-term priority. Second, the ThermoML Archive is actively growing through partnerships with several journals, and new experimental measurements published in these journals are critically examined by the TRC and included in the archive. Finally, the files in the ThermoML Archive are portable and machine readable via a formal XML schema, allowing facile access to hundreds of thousands of measurements. Numerous additional physical properties contained in ThermoML—including activity coefficients, diffusion constants, boiling-point temperatures, critical pressures and densities, coefficients of expansion, speed-of-sound measurements, viscosities, excess molar enthalpies, heat capacities, and volumes—

for neat phases and mixtures represent a rich dataset of high utility for forcefield validation and parameterization.

5 Conclusions

High quality, machine-readable datasets of physicochemical measurements will be required for the construction of next-generation small molecule forcefields. Here we have discussed the NIST/TRC ThermoML archive as a growing source of physicochemical measurements that may be useful for the forcefield community. From the NIST/TRC ThermoML archive, we selected a dataset of 246 ambient, neat liquid systems for which both densities and static dielectric constants are available. Using this dataset, we benchmarked GAFF/AM1-BCC, one of the most popular small molecule forcefields. We find systematic biases in densities and particularly static dielectric constants. Element-based empirical polarizability models are able to account for much of the systematic differences between GAFF/AM1-BCC and experiment. Non-polarizable forcefields may show unacceptable biases when predicting transfer and binding properties of non-polar environments such as binding cavities or membranes.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Patrick B. Grinaway (MSKCC), Vijay S. Pande (Stanford University), Lee-Ping Wang (Stanford University), Peter Eastman (Stanford University), Robert McGibbon (Stanford University), Jason Swails (Rutgers University), David L. Mobley (University of California, Irvine), Michael R. Shirts (University of Virginia), William C. Swope (IBM), Julia E. Rice (IBM), Hans Horn (IBM), Jed W. Pitera (IBM), and members of Chodera lab for helpful discussions. Support for JMB was provided by the Tri-Institutional Training Program in Computational Biology and Medicine (via NIH training grant 1T32GM083937). KAB was supported in part by Starr Foundation grant I8-A8-058. JDC and KAB acknowledge partial support from NIH grant P30 CA008748. KAB, JMB, ASR, and JDC acknowledge the generous support of this research by the Sloan Kettering Institute. The authors gratefully acknowledge OpenEye Scientific for generously providing their toolkit for use in noncommercial projects that generate results for the public domain.

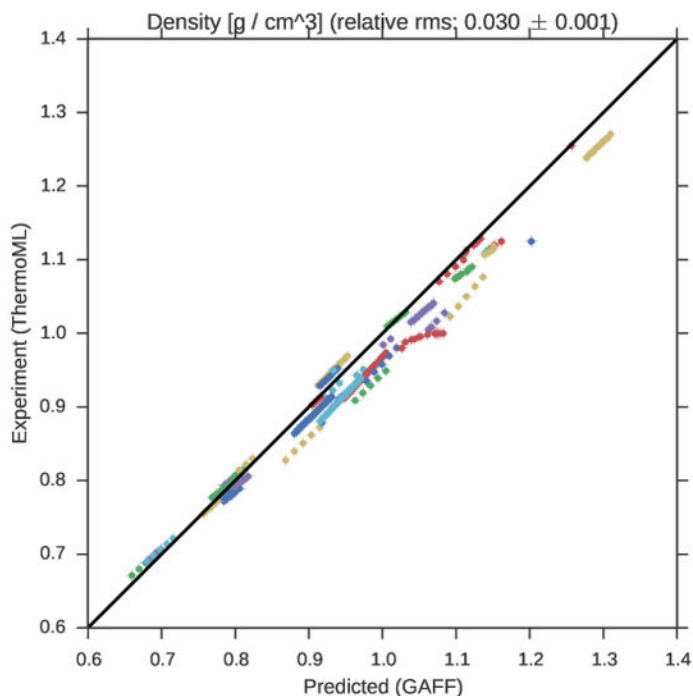
References

1. Salomon-Ferrer R, Gotz AW, Poole D, Le Grand S, Walker RC. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 2. Explicit Solvent Particle Mesh Ewald. *J Chem Theory Comp.* 2013; 9:3878–3888.
2. Lindorff-Larsen K, Maragakis P, Piana S, Eastwood M, Dror R, Shaw D. Systematic Validation of Protein Force Fields against Experimental Data. *PloS one.* 2012; 7:e32131. [PubMed: 22384157]
3. Beauchamp K, Lin Y, Das R, Pande V. Are Protein Force Fields Getting Better? A Systematic Benchmark on 524 Diverse NMR Measurements. *J Chem Theory Comput.* 2012; 8:1409–1414. [PubMed: 22754404]
4. Best R, Buchete N, Hummer G. Are Current Molecular Dynamics Force Fields too Helical? *Biophys J.* 2008; 95:L07–L09. [PubMed: 18456823]
5. Li D-W, Bruschweiler R. Iterative Optimization of Molecular Mechanics Force Fields from NMR Data of Full Length Proteins. *J Chem Theory Comput.* 2011; 7:1773–1782. [PubMed: 26596440]
6. Best RB, Zhu X, Shim J, Lopes PE, Mittal J, Feig M, MacKerell AD. Optimization of the Additive CHARMM All-Atom Protein Force Field Targeting Improved Sampling of the Backbone ϕ , ψ and Side-Chain χ_1 and χ_2 Dihedral Angles. *J Chem Theory Comput.* 2012; 8:3257–3273. [PubMed: 23341755]

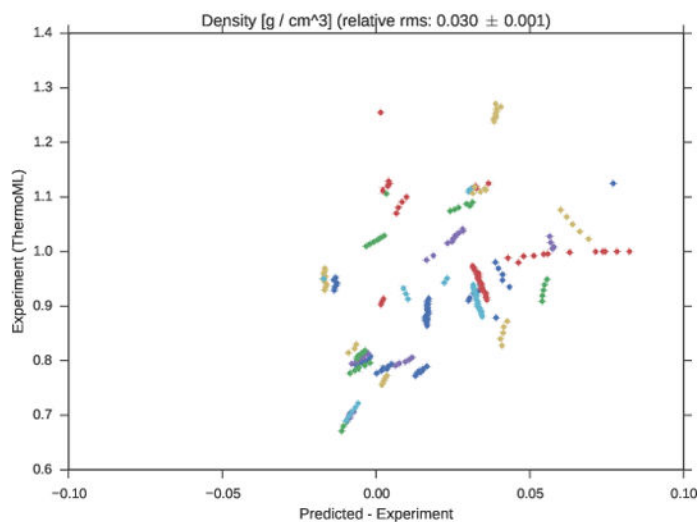
7. Lindorff-Larsen K, Piana S, Palmo K, Maragakis P, Klepeis J, Dror R, Shaw D. Improved Side-Chain Torsion Potentials for the Amber ff99SB Protein Force Field. *Proteins: Struct, Funct Bioinf.* 2010; 78:1950–1958.
8. Lindorff-Larsen K, Piana S, Dror R, Shaw D. How Fast-Folding Proteins Fold. *Science.* 2011; 334:517–520. [PubMed: 22034434]
9. Ensign D, Kasson P, Pande V. Heterogeneity Even at the Speed Limit of Folding: Large-scale Molecular Dynamics Study of a Fast-folding Variant of the Villin Headpiece. *J Mol Biol.* 2007; 374:806–816. [PubMed: 17950314]
10. Voelz V, Bowman G, Beauchamp K, Pande V. Molecular Simulation of ab Initio Protein Folding for a Millisecond Folder NTL9 (1–39). *J Am Chem Soc.* 2010; 132:1526–1528. [PubMed: 20070076]
11. Horn H, Swope W, Pitara J, Madura J, Dick T, Hura G, Head-Gordon T. Development of an Improved Four-site Water Model for Biomolecular Simulations: TIP4P-Ew. *J Chem Phys.* 2004; 120:9665–9678. [PubMed: 15267980]
12. Caleman C, van Maaren PJ, Hong M, Hub JS, Costa LT, van der Spoel D. Force Field Benchmark of Organic Liquids: Density, Enthalpy of Vaporization, Heat Capacities, Surface Tension, Isothermal Compressibility, Volumetric Expansion Coefficient, and Dielectric Constant. *J Chem Theory Comp.* 2011; 8:61–74.
13. Fischer NM, van Maaren PJ, Ditz JC, Yildirim A, van der Spoel D. Properties of Organic Liquids when Simulated with Long-Range Lennard-Jones interactions. *J Chem Theory Comp.* 2015; 11:2938–2944.
14. Zhang J, Tuguldur B, van der Spoel D. Force Field Benchmark of Organic Liquids II: Gibbs Energy of Solvation. *J Chem Inf Mod.* 2015; 55:1192–1201.
15. Fennell CJ, Wymer KL, Mobley DL. 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. [PubMed: 24702668]
16. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The Protein Data Bank. *Nucleic Acids Res.* 2000; 28:235–242. [PubMed: 10592235]
17. Mobley, DL. Experimental and Calculated Small Molecule Hydration Free Energies. UC Irvine: Department of Pharmaceutical Sciences, UCI; Retrieved from: <http://www.escholarship.org/uc/item/6sd403pz>
18. Ulrich E, Akutsu H, Doreleijers J, Harano Y, Ioannidis Y, Lin J, Livny M, Mading S, Maziuk D, Miller Z. BioMagResBank. *Nucleic Acids Res.* 2008; 36:D402–D408. [PubMed: 17984079]
19. Pontolillo, J.; Eganhouse, RP. The Search for Reliable Aqueous Solubility (S_w) and Octanol-Water Partition Coefficient (K_{ow}) Data for Hydrophobic Organic Compounds: DDT and DDE as a Case Study. U.S. Geological Survey; Reston, Virginia: 2001.
20. Frenkel M, Chirico RD, Diky VV, Dong Q, Frenkel S, Francois PR, Embry DL, Teague TL, Marsh KN, Wilhoit RC. ThermoML an XML-based Approach for Storage and Exchange of Experimental and Critically Evaluated Thermophysical and Thermochemical Property Data. 1. Experimental data. *J Chem Eng Data.* 2003; 48:2–13.
21. Frenkel M, Chirico RD, Diky V, Dong Q, Marsh KN, Dymond JH, Wake-ham WA, Stein SE, Königsberger E, Goodwin AR. 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.
22. Chirico RD, Frenkel M, Magee JW, Diky V, Muzny CD, Kazakov AF, Kroenlein K, Abdulagatov I, Hardin GR, Acree WE Jr. Improvement of Quality in Publication of Experimental Thermophysical Property Data: Challenges, Assessment Tools, Global Implementation, and Online Support. *J Chem Eng Data.* 2013; 58:2699–2716.
23. Wang J, Wolf RM, Caldwell JW, Kollman PA, Case DA. Development and Testing of a General AMBER Force Field. *J Comput Chem.* 2004; 25:1157–1174. [PubMed: 15116359]
24. Wang J, Wang W, Kollman PA, Case DA. Automatic Atom Type and Bond Type Perception in Molecular Mechanics Calculations. *J Mol Graph Model.* 2006; 25:247–260. [PubMed: 16458552]
25. Jakalian A, Bush BL, Jack DB, Bayly CI. Fast, Efficient Generation of High-quality Atomic Charges. AM1-BCC Model: I. Method. *J Comput Chem.* 2000; 21:132–146.

26. Jakalian A, Jack DB, Bayly CI. Fast, Efficient Generation of High-quality Atomic Charges. AM1-BCC Model: II. Parameterization and validation. *J Comput Chem.* 2002; 23:1623–1641. [PubMed: 12395429]
27. McKinney W. Data Structures for Statistical Computing in Python. Proceedings of the 9th Python in Science Conference. 2010:51–56.
28. OpenEye Toolkits Version 2014. 2014. <http://www.eyesopen.com>
29. Martínez L, Andrade R, Birgin EG, Martínez JM. Packmol: A Package for Building Initial Configurations for Molecular Dynamics Simulations. *J Comp Chem.* 2009; 30:2157–2164. [PubMed: 19229944]
30. Velez-Vega C, McKay DJ, Aravamuthan V, Pearlstein R, Duca JS. Time-Averaged Distributions of Solute and Solvent Motions: Exploring Proton Wires of GFP and PFM2DH. *J Chem Inf Mod.* 2014; 54:3344–3361.
31. Hawkins PC, Nicholls A. Conformer Generation with OMEGA: Learning from the Data Set and the Analysis of Failures. *J Chem Inf Mod.* 2012; 52:2919–2936.
32. Case, D.; Babin, V.; Berryman, J.; Betz, R.; Cai, Q.; Cerutti, D.; Cheatham, T., III; Darden, T.; Duke, R.; Gohlke, H., et al. AMBER 14. University of California; San Francisco, CA: 2014.
33. Eastman P, Friedrichs MS, Chodera JD, Radmer RJ, Bruns CM, Ku JP, Beauchamp KA, Lane TJ, Wang L-P, Shukla D, et al. OpenMM 4: A Reusable, Extensible, Hardware Independent Library for High Performance Molecular Simulation. *J Chem Theory Comput.* 2012; 9:461–469. [PubMed: 23316124]
34. McGibbon, RT.; Beauchamp, KA.; Schwantes, CR.; Wang, LP.; Hernández, CX.; Harrigan, MP.; Lane, TJ.; Swails, JM.; Pande, VS. MDTraj: a Modern, Open Library for the Analysis of Molecular Dynamics Trajectories. 2014. <http://biorxiv.org/content/early/2014/09/09/008896.short>, accessed Aug. 27, 2015
35. Liu DC, Nocedal J. On the Limited Memory BFGS Method for Large Scale Optimization. *Mathematical programming.* 1989; 45:503–528.
36. Izaguirre, JA.; Sweet, CR.; Pande, VS. Multiscale Dynamics of Macromolecules using Normal Mode Langevin. Pacific Symposium on Biocomputing; Waimee, HI: 2010.
37. Darden T, York D, Pedersen L. Particle Mesh Ewald: An N Log (N) Method for Ewald Sums in Large Systems. *J Chem Phys.* 1993; 98:10089–10092.
38. Chodera, JD. A Simple Method for Automated Equilibration Detection in Molecular Simulations. 2015. <http://biorxiv.org/content/early/2015/07/04/021659>, accessed Aug. 15, 2015
39. Shirts MR, Chodera JD. Statistically Optimal Analysis of Samples from Multiple Equilibrium States. *J Chem Phys.* 2008; 129:124105. [PubMed: 19045004]
40. Chodera JD, Singhal N, Pande VS, Dill KA, Swope WC. Automatic Discovery of Metastable States for the Construction of Markov models of Macromolecular Conformational Dynamics. *J Chem Phys.* 2007; 126:155101. [PubMed: 17461665]
41. Brooks, S.; Gelman, A.; Jones, G.; Meng, XL. Handbook of Markov Chain Monte Carlo. CRC press; Boca Raton, FL: 2011.
42. Neumann M. Dipole Moment Fluctuation Formulas in Computer Simulations of Polar Systems. *Mol Phys.* 1983; 50:841–858.
43. Sheppard, K. ARCH Toolbox for Python. 2015. <http://dx.doi.org/10.5281/zenodo.15681>, GitHub repository: <https://github.com/bashtage/arch>
44. Flyvbjerg H, Petersen HG. Error Estimates on Averages of Correlated Data. *J Chem Phys.* 1989; 91:461.
45. Haider N. Functionality Pattern Matching as an Efficient Complementary Structure/Reaction Search Tool: an Open-source Approach. *Molecules.* 2010; 15:5079–5092. [PubMed: 20714286]
46. Jorgensen WL, Chandrasekhar J, Madura JD, Impey RW, Klein ML. Comparison of Simple Potential Functions for Simulating Liquid Water. *J Chem Phys.* 1983; 79:926–935.
47. Jorgensen WL, Madura JD, Swenson CJ. Optimized Intermolecular Potential Functions for Liquid Hydrocarbons. *J Am Chem Soc.* 1984; 106:6638–6646.
48. Bosque R, Sales J. Polarizabilities of Solvents from the Chemical Composition. *J Chem Inf and Comp Sci.* 2002; 42:1154–1163.

49. Wang L-P, Martínez TJ, Pande VS. Building Force Fields-an Automatic, Systematic and Reproducible Approach. *J Phys Chem Let.* 2014; 5:1885–1891. [PubMed: 26273869]
50. Fennell CJ, Li L, Dill KA. Simple Liquid Models with Corrected Dielectric Constants. *J Phys Chem B.* 2012; 116:6936–6944. [PubMed: 22397577]
51. Leontyev IV, Stuchebrukhov AA. Polarizable Molecular Interactions in Condensed Phase and their Equivalent Nonpolarizable Models. *J Chem Phys.* 2014; 141:014103. [PubMed: 25005273]
52. D’Aprano A, Donato ID. Dielectric Polarization and Polarizability of 1-pentanol+ n-octane Mixtures from Static Dielectric Constant and Refractive Index Data at 0, 25 and 45 C. *J Sol Chem.* 1990; 19:883–892.
53. Haynes, WM. *CRC handbook of chemistry and physics.* CRC Press; Boca Raton, FL: 2011.
54. Mobley, DL. Experimental and Calculated Small Molecule Hydration Free Energies. UC Irvine: Department of Pharmaceutical Sciences, UCI; Retrieved from: <https://github.com/choderalab/FreeSolv>
55. Newman J, Fazio VJ, Caradoc-Davies TT, Branson K, Peat TS. Practical Aspects of the SAMPL Challenge: Providing an Extensive Experimental Data Set for the Modeling Community. *J Biomol Screening.* 2009; 14:1245–1250.
56. Truchon J-F, Nicholl’s A, Grant JA, Iftimie RI, Roux B, Bayly CI. Using Electronic Polarization from the Internal Continuum (EPIC) for Intermolecular Interactions. *J Comp Chem.* 2010; 31:811–824. [PubMed: 19598266]
57. Truchon J-F, Nicholls A, Roux B, Iftimie RI, Bayly CI. Integrated Continuum Dielectric Approaches to Treat Molecular Polarizability and the Condensed Phase: Refractive Index and Implicit Solvation. *J Chem Theory Comp.* 2009; 5:1785–1802.
58. Truchon J-F, Nicholls A, Iftimie RI, Roux B, Bayly CI. Accurate Molecular Polarizabilities Based on Continuum Electrostatics. *J Chem Theory Comp.* 2008; 4:1480–1493.
59. Ponder J, Wu C, Ren P, Pande V, Chodera J, Schnieders M, Haque I, Mobley D, Lambrecht D, DiStasio R Jr, et al. Current Status of the AMOEBA Polarizable Force Field. *J Phys Chem B.* 2010; 114:2549–2564. [PubMed: 20136072]
60. Ren P, Ponder JW. Temperature and Pressure Dependence of the AMOEBA Water Model. *J Phys Chem B.* 2004; 108:13427–13437.
61. Lamoureux G, Roux B. Modeling Induced Polarization with Classical Drude Oscillators: Theory and Molecular Dynamics Simulation Algorithm. *J Chem Phys.* 2003; 119:3025–3039.
62. Anisimov VM, Lamoureux G, Vorobyov IV, Huang N, Roux B, MacKerell AD. Determination of Electrostatic Parameters for a Polarizable Force Field Based on the Classical Drude Oscillator. *J Chem Theory Comp.* 2005; 1:153–168.
63. Wang L-P, Head-Gordon TL, Ponder JW, Ren P, Chodera JD, East-man PK, Martínez TJ, Pande VS. Systematic Improvement of a Classical Molecular Model of Water. *J Phys Chem B.* 2013; 117:9956–9972. [PubMed: 23750713]
64. Chirico RD, Frenkel M, Diky VV, Marsh KN, Wilhoit RC. ThermoML An XML-Based Approach for Storage and Exchange of Experimental and Critically Evaluated Thermophysical and Thermochemical Property Data. 2. Uncertainties. *J Chem Eng Data.* 2003; 48:1344–1359.



(a)



(b)

Figure 1. Comparison of liquid densities between experiment and simulation

(a). Liquid density measurements extracted from ThermoML are compared against densities predicted using the GAFF/AM1-BCC small molecule fixed-charge forcefield. Color groupings represent identical chemical species, although the color map repeats itself due to the large (45) number of unique compounds. Plots of density versus temperature grouped by chemical species are available in Fig. S4. Simulation error bars represent one standard error of the mean, with the number of effective (uncorrelated) samples estimated using pymbar. Experimental error bars indicate the standard deviation between independently reported

measurements, when available, or author-reported standard deviations in ThermoML entries; for some measurements, neither uncertainty nor estimate is available. See Fig. S2 for further discussion of error. (b). The same plot, but with the residual (predicted minus experiment) on the x axis. Note that the error bars are all smaller than the symbols.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

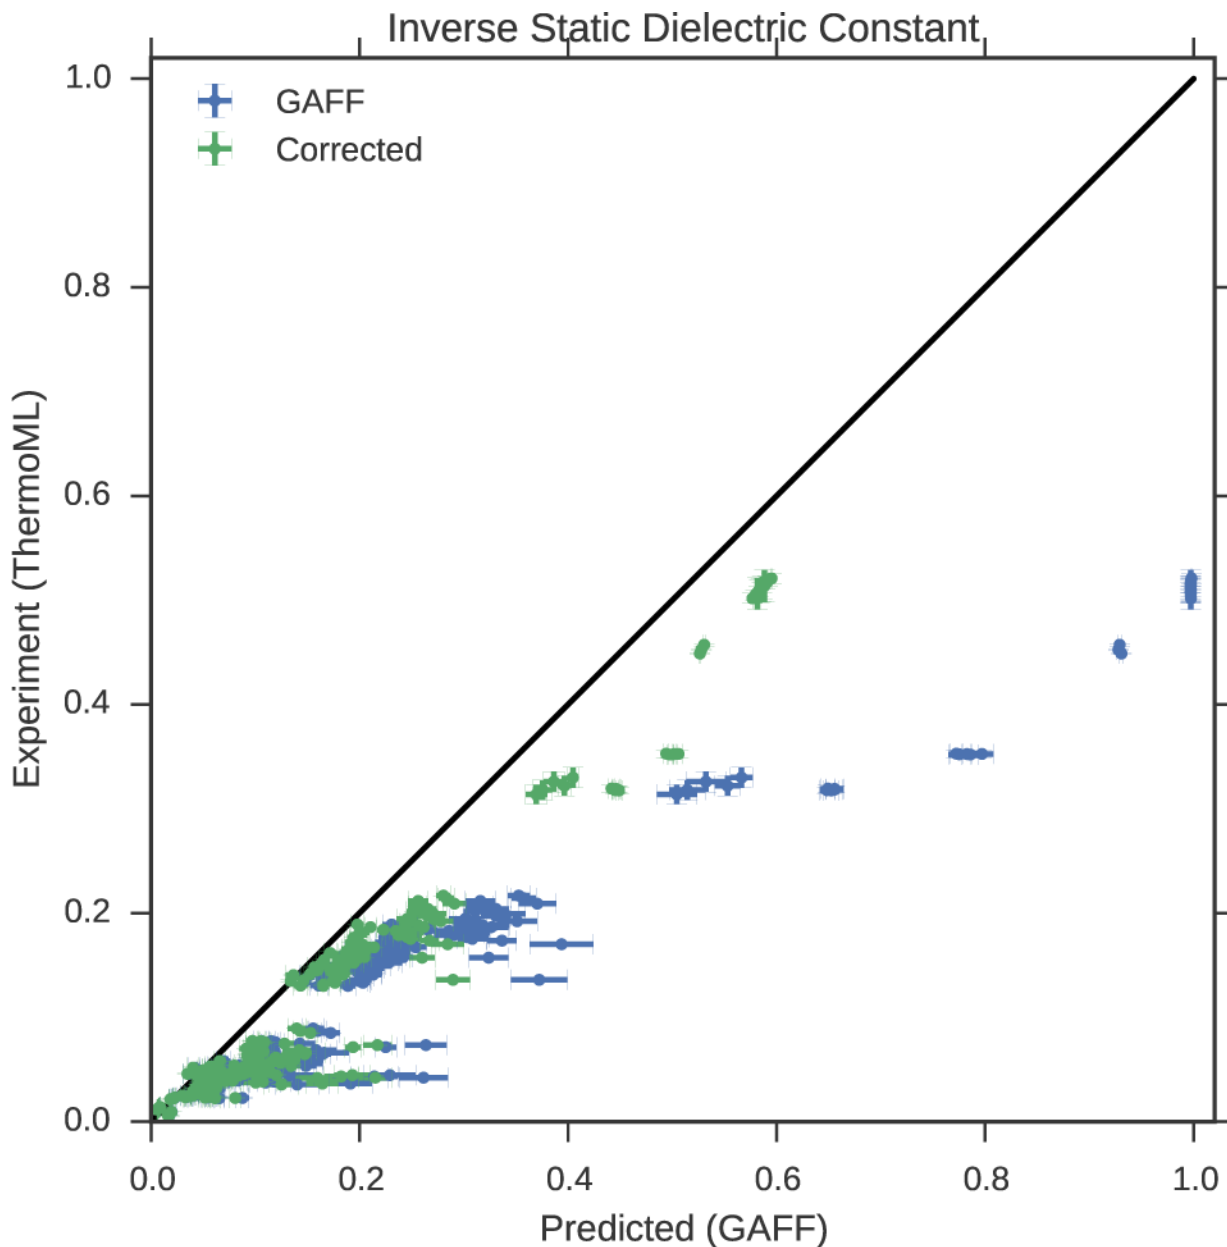


Figure 2. Measured (ThermoML) versus predicted (GAFF/AM1-BCC) inverse static dielectrics (a). Simulation error bars represent one standard error of the mean. Experimental error bars indicate the larger of standard deviation between independently reported measurements and the authors reported standard deviations; for some measurements, neither uncertainty estimate is available. See Fig. S3 for further discussion of error. See Section 3.2.2 for explanation of why inverse dielectric constant (rather than dielectric constant) is plotted. For nonpolar liquids, it is clear that the forcefield predicts electrostatic interactions that are substantially biased by missing polarizability. Plots of dielectric constant versus temperature grouped by chemical species are available in Fig. S5.




Figure 3. Typical experimental static dielectric constants of some nonpolar compounds
 (a). Measured static dielectric constants of various nonpolar or symmetric molecules.^{52,53} Fixed-charge forcefields give $\epsilon \approx 1$ for each species; for example, we calculated $\epsilon = 1.0030 \pm 0.0002$ for octane. (b). A congenetic series of chloro-substituted methanes have static dielectric constants between 2 and 13. Reported dielectric constants are at near-ambient temperatures.

Table 1

Successive filtration of the ThermoML Archive. A set of successive filters were applied to all measurements in the ThermoML Archive that contained either mass density or static dielectric constant measurements. Each column reports the number of measurements remaining after successive application of the corresponding filtration step. The 246 final measurements correspond to 45 unique molecules measured at several temperature conditions.

Filter step	Number of measurements remaining	
	Mass density	Static dielectric
1. Single Component	136212	1651
2. Druglike Elements	125953	1651
3. Heavy Atoms	71595	1569
4. Temperature	38821	964
5. Pressure	14103	461
6. Liquid state	14033	461
7. Aggregate T, P	3592	432
8. Density+Dielectric	246	246

Table 2

Functional groups present in filtered dataset. The filtered ThermoML dataset contained 246 distinct (molecule, temperature, pressure) conditions, spanning 45 unique compounds. The functional groups represented in these compounds (as identified by the program checkmol v0.5⁴⁵) is summarized here.

Functional Group	Occurrences
1,2-aminoalcohol	4
1,2-diol	3
alkene	3
aromatic compound	1
carbonic acid diester	2
carboxylic acid ester	4
dialkyl ether	7
heterocyclic compound	3
ketone	3
lactone	1
primary alcohol	19
primary aliphatic amine (alkylamine)	2
primary amine	2
secondary alcohol	4
secondary aliphatic amine (dialkylamine)	2
secondary aliphatic/aromatic amine (alkylarylamine)	1
secondary amine	3
sulfone	1
sulfoxide	1
tertiary aliphatic amine (trialkylamine)	3
tertiary amine	3