

# Machine Learning Directed Optimization of Classical Molecular Modeling Force Fields

Bridgette J. Befort,<sup>‡</sup> Ryan S. DeFever,<sup>‡</sup> Garrett M. Tow, Alexander W. Dowling, and Edward J. Maginn\*



Cite This: <https://doi.org/10.1021/acs.jcim.1c00448>



Read Online

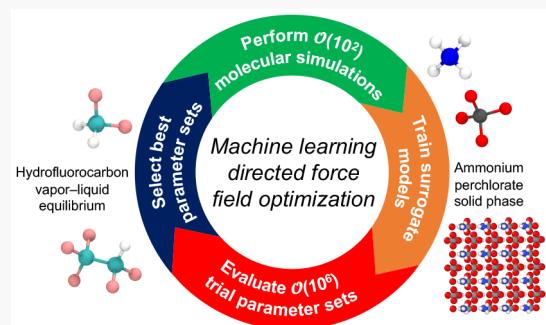
ACCESS |

Metrics & More

Article Recommendations

Supporting Information

**ABSTRACT:** Accurate force fields are necessary for predictive molecular simulations. However, developing force fields that accurately reproduce experimental properties is challenging. Here, we present a machine learning directed, multiobjective optimization workflow for force field parametrization that evaluates millions of prospective force field parameter sets while requiring only a small fraction of them to be tested with molecular simulations. We demonstrate the generality of the approach and identify multiple low-error parameter sets for two distinct test cases: simulations of hydrofluorocarbon (HFC) vapor–liquid equilibrium (VLE) and an ammonium perchlorate (AP) crystal phase. We discuss the challenges and implications of our force field optimization workflow.



## INTRODUCTION

Molecular modeling and simulation use computational methods to describe the behavior of matter at the atomistic or molecular level.<sup>1</sup> The veracity and predictive capability of molecular simulations depend critically on the accuracy of the atomic-level interaction energies, and whether the appropriate time- and length-scales are properly sampled. On one hand is a class of techniques broadly termed as *ab initio* or first-principles methods, where atomic interactions are determined from highly accurate quantum chemical methods.<sup>2</sup> Though there are applications that necessitate these methods, *ab initio* energies are computationally expensive to obtain, such that quantum chemical methods are limited to relatively small systems and short time scales. On the other hand, classical molecular simulations represent the atomic interaction energies with an analytical function (a “force field”) that can be evaluated much more rapidly than *ab initio* energies enabling simulations of much larger systems and longer time scales than is possible with *ab initio* techniques. If force fields are highly accurate, classical molecular simulations have been shown to give accurate property predictions in several fields including protein structure refinement,<sup>3</sup> drug discovery,<sup>4</sup> and energy storage.<sup>5</sup>

**Developing Accurate Force Fields is Difficult.** There are two fundamentally different approaches to developing and improving force fields: bottom-up approaches, wherein parameters are calibrated so the model reproduces the results (e.g., forces, energies, and dipoles) of more expensive and accurate methods (i.e., quantum calculations),<sup>6</sup> and top-down approaches, wherein parameters are calibrated so the model matches experimental results.<sup>7</sup> Emerging bottom-up approaches use machine learning (ML) to parametrize force

fields with black-box potential energy functions.<sup>8,9</sup> Though these so-called ML force fields<sup>10,11</sup> have proven successful for an increasing number of systems, the black-box nature of the potential energy function makes the models physically uninterpretable and hinders model transferability beyond the specific training conditions. Developing accurate and transferable force fields with analytical functional forms is a difficult and laborious endeavor.<sup>12</sup> Significant efforts spanning several decades have resulted in several “off-the-shelf” force fields that describe large swaths of condensed matter chemical space.<sup>13–16</sup> These are most commonly “Class I” force fields that consist of harmonic or sinusoidal intramolecular terms that describe bonded interactions, atomic partial charges that represent electrostatic interactions, and nonbonded repulsion–dispersion terms. Unfortunately, these off-the-shelf force fields can yield poor property predictions, even for relatively common compounds, particularly when they are applied in circumstances beyond the systems and conditions for which they were parametrized.<sup>17</sup> However, since they are well-known and the parameter sets are widely distributed, these force fields are used in many molecular simulation studies.

For decades, force field development and optimization has been an active area of research. Several methods and tools have been developed to derive bonded intramolecular parameters and partial charges in a bottom-up fashion from quantum

Received: April 19, 2021

calculations, provided that the desired classical functional form has been selected. Common approaches include gradient-based techniques, evolutionary algorithms, or even analytical solutions.<sup>18–25</sup> These methods work well because the relevant quantities can be computed to a high degree of accuracy with quantum calculations and evaluating a prospective force field parameter set is computationally trivial. However, optimizing the repulsion–dispersion parameters that are largely responsible for many macroscopic thermodynamic properties (e.g., density, enthalpy of vaporization, vapor pressure, etc.) is more challenging. Since these parameters can be difficult to derive from quantum calculations without special methods,<sup>26</sup> top-down parametrization is often necessary. Yet screening thousands of prospective parameter sets is computationally expensive due to the need for sufficiently long simulations to accurately compute the relevant experimental properties. Even for relatively simple properties, a single simulation can require hours-to-days of computation time.

It is often desirable to parametrize a force field to reproduce multiple physical properties. A rigorous way to calibrate force fields with multiple properties simultaneously is to use multiobjective optimization,<sup>27–31</sup> which can exacerbate the computational burden by an order of magnitude or more. In multiobjective optimization, a solution is Pareto optimal if it is not possible to improve one objective without sacrificing another objective.<sup>32</sup> One approach is to weight each objective and re-solve the optimization problem for many different weights to identify Pareto optimal solutions.<sup>33</sup> Thus, computing a set of Pareto optimal solutions is often at least an order of magnitude more computationally expensive than single objective optimization. With much less computational effort, a finite set of candidate solutions can be classified into two groups: the nondominated set, which comprises the solutions for which no other solution in the set offers improvement in any one objective without degrading performance in another objective, and the dominated set, comprising the solutions for which another solution offers improved performance in one or more objectives without degrading the performance in any other objective. By definition, all points in the Pareto set are nondominated; the nondominated set is an easy to compute approximation of the Pareto set.

Given the challenges associated with top-down optimization of the repulsion–dispersion parameters, there are fewer methods and packages available<sup>34,35</sup> compared to intramolecular parameters and partial charge optimization. Much more frequently, attempts to improve these parameters involve *ad hoc* hand-tuning,<sup>36,37</sup> which is arbitrary and often limited to a few interaction parameters or a scaling thereof, as larger searches quickly become intractable.<sup>38</sup> Instead of performing multiobjective optimization, the more common approach is to use *ad hoc* weights to combine multiple calibration objectives into a single cost function.<sup>28,34,35</sup> However, this approach only finds a single Pareto optimal trade-off between the calibration objectives.

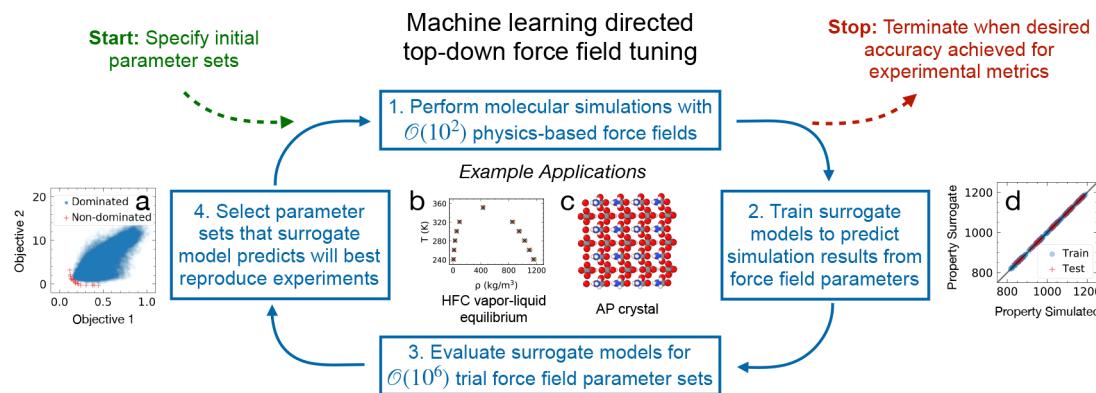
**Machine Learning Directed Optimization Makes Force Field Calibration More Computationally Tractable.** The core challenges of optimizing the repulsion–dispersion parameters can be solved with a computationally inexpensive mapping between the desired physical properties and force field parameters. For certain cases, these mappings can be constructed with statistical mechanics,<sup>39,40</sup> but this approach likely cannot be generalized to arbitrary systems. Alternatively, ML can be used to approximate the relevant

mapping. For example, surrogate-assisted optimization (also known as black-box or derivative-free optimization) uses computationally inexpensive surrogate model evaluations to emulate the outputs of a complex computer simulation, e.g., computational fluid dynamics, finite element analysis, or molecular simulations. Several different types of surrogate models have been successfully applied to molecular simulations for uncertainty quantification<sup>41,42</sup> and force field parametrization.<sup>35,43–45</sup> Linear regression response surface models were used to predict the optimal combination of scaling factors for the charge and Lennard-Jones (LJ) parameters of General AMBER force field (GAFF) to reproduce four properties of organic liquid electrolytes. While easy to implement and moderately successful at improving the force field's accuracy for some of the properties, this method was limited by the choice of statistically significant parameters in the response surface.<sup>46</sup> For some thermodynamic properties, reweighting methods are an effective tool to test a large number of parameters without performing additional simulations,<sup>44,47,48</sup> but care must be taken to ensure good phase space overlap between the sampled and reweighted ensembles.<sup>44</sup> Gaussian process regression (GPR) is a popular nonparametric surrogate model that smoothly interpolates between training data. Some applications of GPR in molecular simulations include ML force fields<sup>49–51</sup> and property prediction.<sup>52</sup> In Bayesian optimization, which is a special case of surrogate-assisted optimization, the uncertainty estimates from GPR (or a similar model) are directly used to balance exploration and exploitation. Recent work demonstrates Bayesian optimization can efficiently calibrate force field parameters in coarse-grained models.<sup>53–55</sup> Moreover, computationally inexpensive surrogate models can enable multiobjective optimization algorithms that go beyond *ad hoc* weighting<sup>32</sup> to systematically explore trade-offs when calibrating multiple physical properties.

Here, we demonstrate a new multiobjective surrogate-assisted optimization framework that uses GPRs and support vector machine (SVM) classifiers to improve existing all-atom force fields. The proposed strategy enables extremely accurate property calculations while retaining physically motivated and interpretable functional forms. We show that the same general approach successfully optimizes force fields for two systems with very different characteristics and property objectives: hydrofluorocarbon (HFC) vapor–liquid equilibrium (VLE) and solid ammonium perchlorate (AP) crystal structure. Our results highlight the versatility of surrogate-assisted optimization approaches for top-down parametrization of all-atom force fields in a wide range of domains. The remainder of the manuscript proceeds as follows: we outline the method and provide technical details in the **Methodology** section, demonstrate the approach for the two case studies in the **Results**, discuss the challenges and implications of the method in the **Discussion** section, and provide concluding remarks in the **Conclusions**.

## ■ METHODOLOGY

**Machine Learning Directed Force Field Optimization Workflow.** An overview of our force field optimization workflow is provided first with a more technical description given in the following subsections. Our strategy in this work is to optimize LJ repulsion–dispersion parameters, which are among the most difficult to calculate from *ab initio* methods.<sup>56</sup> Intramolecular parameters and partial charges, which usually can be reliably and inexpensively determined from bottom-up



**Figure 1.** Overview of the proposed machine learning directed force field optimization procedure. The workflow tests  $\mathcal{O}(10^6)$  sets of force field parameters for every  $\mathcal{O}(10^2)$  molecular simulations. The four main steps are described in the numbered boxes. (a) Difference between dominated and nondominated solutions for an example where the goal is to minimize two objectives. (b and c) Two example applications. (d) Example of how the surrogate models accurately predict the outcomes of molecular simulations.

*ab initio*-based methods, were determined from existing force fields. We stress, however, that this method can be applied to calibrate any force field parameters.

Our force field optimization workflow is shown schematically in Figure 1. First, domain knowledge is used to specify physically reasonable bounds on the search space for the parameters that are being optimized. Next,  $\mathcal{O}(10^2)$  initial parameter sets are generated via space-filling Latin hypercube sampling (LHS). Molecular simulations are performed with each parameter set (Figure 1, box 1), and the physical properties of interest are computed from the simulations. These results are used to train surrogate models (box 2, panel d) that predict the simulation results directly from the parameter set, and optionally, the thermodynamic state point, e.g.,  $T$  and  $p$ . Additional examples of surrogate model accuracy can be found in SI Figures S1 and S2. The surrogate model is then used to predict the molecular simulation results for a very large number,  $\mathcal{O}(10^6)$ , of candidate parameter sets, once again generated with LHS (box 3). The  $\mathcal{O}(10^2)$  most promising parameter sets are identified via user-selected system-specific metrics including error thresholds, separation in parameter space, and nondominated status, from the  $\mathcal{O}(10^6)$  candidate sets evaluated with the surrogate models (box 4). In multiobjective optimization, the set of nondominated points includes all parameter sets that are not simultaneously outperformed in every dimension by any other parameter set (Figure 1a).<sup>32</sup> Finally, the most promising parameter sets are used to initialize the next iteration of molecular simulations (box 1). The process is repeated until parameter sets are generated that provide the desired accuracy for the experimental properties of interest.

The workflow uses a combination of machine learning-based surrogate models and physics-based molecular simulations to quickly optimize force field parameters for a specific system. Physically motivated potential energy functional forms that have proven successful over decades are retained. Whereas the molecular simulations require hours-to-days to compute experimentally measurable properties arising from a single set of force field parameters, the surrogate models can evaluate millions of parameter sets in minutes-to-hours. This means that once the surrogate models have been trained to predict the results of the molecular simulations, they enable an exhaustive search of large parameter spaces that would require

$\mathcal{O}(10^7\text{--}10^9)$  CPU-hours with molecular simulations. We emphasize that although the surrogate models are used to screen millions of candidate parameter sets, all of the promising candidate parameter sets are ultimately tested with physics-based molecular simulations. The role of machine learning is only to act as a surrogate for physics-based simulations, enabling the parameter search through an otherwise intractable space. The iterative procedure allows the surrogate models to improve as additional training data is collected with each iteration. The original molecular simulations are dispersed across the entire parameter space, but subsequent iterations are focused on the smaller regions of parameter space that are predicted to yield good parameter sets, enabling the surrogate models to improve in the most important regions of parameter space. The theory and technical details of each step in Figure 1 are presented below. Methodological details specific to the HFC and AP examples are reported in the Hydrofluorocarbon and Ammonium Perchlorate Case Study sections, respectively.

**Problem Setup.** The interaction potential is taken as a classical molecular mechanics force field,  $U(\mathbf{r}) = f(\mathbf{r}, \boldsymbol{\zeta})$ , where  $U$  is the potential energy,  $\mathbf{r} \in \Gamma$  is the vector of position coordinates within configuration space  $\Gamma$ ,  $f$  is the functional form for the potential energy, and  $\boldsymbol{\zeta} = \zeta_1, \zeta_2, \dots, \zeta_N$  are the parameters of  $f$  that define the intra- and intermolecular interactions between different types of particles. Molecular simulations can be used to compute  $M$  structural, thermodynamic, or dynamic properties,  $\mathbf{y}^{\text{sim}} = y_1^{\text{sim}}, y_2^{\text{sim}}, \dots, y_M^{\text{sim}}$ , from  $U(\mathbf{r})$ . Depending upon the quality of  $U(\mathbf{r})$ ,  $\mathbf{y}^{\text{sim}}$  may or may not be close to the experimental values,  $\mathbf{y}^{\text{exp}}$ . The goal of this work is to refine  $U(\mathbf{r})$  by optimizing  $\mathcal{O}(10^1)$  force field parameters,  $\boldsymbol{\zeta}' \subseteq \boldsymbol{\zeta}$ , such that  $\mathbf{y}^{\text{sim}} \approx \mathbf{y}^{\text{exp}}$  for one or more physical properties of interest. In both case studies presented here, the LJ parameters,  $\sigma$  and  $\epsilon$ , are optimized. Upper and lower bounds for each parameter are selected to span a wide range of physically reasonable values. The initial  $\mathcal{O}(10^2)$  parameter sets are randomly selected to be space-filling within these bounds with LHS.

**Step 1: Perform Molecular Simulations with  $\mathcal{O}(10^2)$  Physics-Based Force Fields.** Molecular simulations are performed for each parameter set with the molecular dynamics (MD) or Monte Carlo (MC) method. For each parameter set,  $\mathbf{y}^{\text{sim}}$  is computed from the simulation output. Simulations may

be performed at multiple thermodynamic conditions (e.g.,  $T$  and  $p$ ) for each parameter set if the experimental data exist. Signac-flow was used to manage the setup and execution of all molecular simulations.<sup>57,58</sup>

**Step 2: Train Surrogate Models to Predict Simulation Results from Force Field Parameters.** Gaussian process (GP) surrogate models are trained to predict  $\mathbf{y}^{\text{sim}}$  as a function of the calibrated parameters  $\boldsymbol{\zeta}'$ . For each property, we train the following:

$$\hat{y}_i^{\text{sim}} = \text{GP}_i(m_i(\boldsymbol{\zeta}'), \text{cov}_i(\boldsymbol{\zeta}', \boldsymbol{\zeta}')) \quad (1)$$

where  $\hat{y}_i^{\text{sim}}$  is the surrogate model prediction of  $y_i^{\text{sim}}$ ,  $\text{GP}_i$  is the GP model for property  $i$ ,  $m_i$  is the mean function, and  $\text{cov}_i$  is the covariance (kernel) function. All GP models were implemented in GPFlow 2.0.0.<sup>59</sup> To improve the accuracy of the GP models in regions of parameter space where  $\mathbf{y}^{\text{sim}} \approx \mathbf{y}^{\text{exp}}$ , we exclude parameter sets that result in extremely poor or unphysical results from the GP training data. We then trained SVM classifiers to predict if a parameter set was unphysical (e.g., simulation fails) so that parameter sets from these regions of parameter space could be excluded when the GP models were used to predict the results of trial parameter sets. All SVM classifiers were implemented in scikit-learn<sup>60</sup> with a radial basis function kernel.

**Step 3: Evaluate Surrogate Models for  $O(10^6)$  Trial Force Field Parameter Sets.** After the GP and SVM models are trained,  $O(10^6)$  trial parameter sets are generated with LHS. For each parameter set, the SVM and GP models are used to calculate  $\hat{\mathbf{y}}^{\text{sim}}$ , the surrogate model estimates of  $\mathbf{y}^{\text{sim}}$ .

**Step 4: Select Parameter Sets that Surrogate Models Predict Will Best Reproduce Experiments.** Parameter sets where the surrogate models predict good agreement with experiment,  $\hat{\mathbf{y}}^{\text{sim}} \approx \mathbf{y}^{\text{exp}}$ , are selected for the next iteration. In some cases we apply an optional distance-based search algorithm (see **SI Methods**) to down-select only parameter sets that are far apart in parameter space.

**Hydrofluorocarbon Case Study.** Force fields were independently developed for two HFCs: difluoromethane (HFC-32) and pentafluoroethane (HFC-125). Two stages of optimization were used for each HFC. The first stage used MD simulations in the  $NpT$  ensemble at 241, 261, 281, 301, and 321 K for HFC-32 and 229, 249, 269, 289, and 309 K for HFC-125. For each temperature, the pressure was set to the experimental<sup>61</sup> saturation pressure. The only property considered during the first stage was the liquid density (LD) ( $\mathbf{y} = \{\rho^l\}$ ). In the second stage of optimization, Gibbs ensemble Monte Carlo (GEMC) was performed. The property objectives were the saturated liquid density, saturated vapor density, vapor pressure, and enthalpy of vaporization, or  $\mathbf{y} = \{\rho_{\text{sat}}^l, \rho_{\text{sat}}^v, P_{\text{vap}}, \Delta H_{\text{vap}}\}$ . Simulations were performed at the same temperatures used for the first stage. Four iterations of the stage 1 optimization were performed for both HFC-32 and HFC-125. Three and five iterations of stage 2 optimization were performed for HFC-32 and HFC-125, respectively.

**Force Field Parameters.** The functional form was taken from GAFF:<sup>15</sup>

$$U(\mathbf{r}) = U^{\text{intra}}(\mathbf{r}) + \sum_i \sum_{j>i} \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}} + \sum_i \sum_{j>i} 4\epsilon_{ij} \left[ \left( \frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left( \frac{\sigma_{ij}}{r_{ij}} \right)^6 \right] \quad (2)$$

where  $U^{\text{intra}}$  contains all the intramolecular terms,  $r_{ij}$  is the distance between atoms  $i$  and  $j$ ,  $q$  is the atomic charge,  $\epsilon_0$  is the permittivity of free space, and  $\sigma_{ij}$  and  $\epsilon_{ij}$  parametrize the LJ potential that describes the repulsion–dispersion interactions between atoms  $i$  and  $j$ . The intramolecular interactions are given by

$$U^{\text{intra}}(\mathbf{r}) = \sum_{\text{bonds}} k_r (r - r_0)^2 + \sum_{\text{angles}} k_\theta (\theta - \theta_0)^2 + \sum_{\text{dihedrals}} \nu_n [1 + \cos(n\phi - \gamma)] \quad (3)$$

where  $r_0$  and  $\theta_0$  are the nominal bond length and angle, respectively,  $k_r$ ,  $k_\theta$ , and  $\nu_n$  are force constants,  $n$  is the multiplicity, and  $\gamma$  is the nominal dihedral angle. The sums are over all bonds, angles, and dihedrals in the system, respectively. The bond, angle, and dihedral parameter for HFC-32 and HFC-125 were taken from GAFF.<sup>15</sup> Partial charges were determined with RESP<sup>18</sup> as implemented in AmberTools 1.4.<sup>62</sup> The quantum electrostatic potential was computed with Gaussian 09<sup>63</sup> with B3LYP/6-311++g(d,p).<sup>64,65</sup> The intramolecular parameters and partial charges are reported in **SI Table S1**.

The force field optimization method was used to determine the like-interaction parameters  $\sigma_{ii}$  and  $\epsilon_{ii}$  for three atom types (C, F, and H) in HFC-32 and five atom types (C1, C2, F1, F2, and H) in HFC-125. This results in 6 parameters that are optimized for HFC-32 and 10 parameters that are optimized for HFC-125. All unlike interaction parameters were computed with Lorentz–Berthelot mixing rules. For HFC-125, C1 is the carbon bonded to one carbon atom, two fluorine atoms, and one hydrogen atom, while C2 is the carbon bonded to one carbon atom and three fluorine atoms, F1 is bonded with to C1, and F2 is bonded with C2. The lower and upper bounds for each parameter were selected per-element ( $\sigma$  in Å,  $\epsilon/k_B$  in K):  $3.0 \leq \sigma_C \leq 4.0$ ,  $2.5 \leq \sigma_F \leq 3.5$ ,  $1.7 \leq \sigma_H \leq 2.7$ ,  $20 \leq \epsilon_C/k_B \leq 60$ ,  $15 \leq \epsilon_F/k_B \leq 40$ ,  $2 \leq \epsilon_H/k_B \leq 10$ . The parameter bounds for each atom type in HFC-32 and HFC-125 are summarized in **SI Tables S2 and S3**, respectively.

**Classifier.** An SVM classifier was trained to predict parameter sets that yielded spontaneous vaporization ( $\rho^l < 500 \text{ kg/m}^3$ ) in MD simulations initiated at liquid density from  $\boldsymbol{\zeta}'$  and  $T$ .

**GP Model.** The GP models predicted the value of a physical property from  $\boldsymbol{\zeta}'$  and  $T$ . The LD iterations used one GP model that predicted  $\rho^l$ . Parameter sets with  $\rho^l < 500 \text{ kg/m}^3$  were excluded from the GP training data. The VLE iterations used one GP model for each property:  $\{\rho_{\text{sat}}^l, \rho_{\text{sat}}^v, P_{\text{vap}}, \Delta H_{\text{vap}}\}$ . All GP models used a radial basis function or Matérn  $\nu = 5/2$  kernel and a linear mean function.<sup>66</sup>

**Selecting Parameter Sets for the Next Iteration.** A new LHS with 1 000 000 (HFC-32) or 500 000 (HFC-125) parameter sets was generated for each iteration. **LD iterations:** Each parameter set was evaluated with the LD SVM classifier at the highest  $T$ . Each parameter set was evaluated with the LD GP model at each  $T$ , and the root-mean-square error (RMSE)

between the GP model prediction and experimental liquid density across all five temperatures was calculated for each parameter set. The 100 lowest RMSE parameter sets that the SVM predicted would remain liquid, and the 100 lowest RMSE parameter sets that the SVM predicted would transform to vapor, were selected for the next iteration. The low-RMSE, predicted-vapor parameter sets were included because they reflect disagreement between the SVM and GP models. After four LD iterations, parameter sets for the VLE-1 iteration were selected from the 800 simulated parameter sets. A distance-based search algorithm (see *SI Methods*) was used to select 25 well-separated parameter sets with  $\text{RMSE} \leq 10 \text{ kg/m}^3$ . **VLE iterations:** Each parameter set from the LHS was evaluated with the LD GP model. Parameter sets predicted to yield LD  $\text{RMSE} > 25 \text{ kg/m}^3$  were discarded. This step was included to make use of the training data generated during the LD iterations since the LD GP model is very accurate after four LD iterations. The remaining parameter sets were evaluated with the four GP models trained to predict VLE properties  $\{\rho_{\text{sat}}^l, \rho_{\text{sat}}^v, P_{\text{vap}}, \Delta H_{\text{vap}}\}$ . The RMSE difference between the GP model predictions and experimental values across all five temperatures was calculated for each property and parameter set. All dominated parameter sets were discarded. A parameter set is dominated if one or more parameter sets performs better than it in all of the considered objective dimensions (e.g., physical properties). The 25 parameter sets selected for the next iteration comprised the top performing parameter set for each physical property and 21 parameter sets selected from the remaining nondominated parameter sets. A distance-based search algorithm identified parameter sets that were well-separated in parameter space.

**MD Simulations.** Simulations of 150 HFC molecules were performed in the  $NpT$  ensemble at the experimental saturation pressure. Initial configurations were generated at  $1000 \text{ kg/m}^3$ . Following a steepest descent energy minimization, systems were equilibrated for 500 ps with the Bussi thermostat<sup>67</sup> and Berendsen barostat<sup>68</sup> with  $\tau_T = 0.1 \text{ ps}$ ,  $\tau_p = 0.5 \text{ ps}$ . The production simulations were 2.5 ns in length with the Bussi thermostat and Parrinello–Rahman barostat<sup>69</sup> with  $\tau_T = 0.5 \text{ ps}$  and  $\tau_p = 1.0 \text{ ps}$ . The final 2.0 ns of the production simulations were used to compute the average density.

The equations of motion were integrated with the leapfrog algorithm<sup>70</sup> and a time step of 1.0 fs. LJ interactions and short-range electrostatics were cut off at 1.0 nm. The particle mesh Ewald method<sup>71</sup> was used to compute long-range electrostatic interactions. Analytical tail corrections to the LJ potential were applied to energy and pressure. All bonds were constrained with the P-LINCS<sup>72</sup> method with the lincs-order and lincs-iter set to 8 and 4, respectively. Simulations were performed with GROMACS 2020.<sup>73</sup>

**MC Simulations.** GEMC simulations were performed with 1000 HFC molecules. The initial liquid box (800 HFC molecules) was generated at the experimental liquid density and pre-equilibrated with a 5000 sweep  $NpT$  MC simulation. The initial vapor box (200 HFC molecules) was randomly generated at the vapor density estimated from the ideal gas law. The combined system was simulated with GEMC. The systems were equilibrated for 10 000 MC sweeps followed by a production GEMC simulation that was 90 000 MC sweeps.

LJ interactions and short-range electrostatics were cut off at 1.2 nm in the liquid box and 2.5 nm in the vapor box. Long-range electrostatics were computed with an Ewald summation with a relative accuracy of  $10^{-5}$ . Analytical tail corrections to

the LJ interactions were applied to energy and pressure. All bonds were fixed at their nominal bond length. Simulations were performed with MoSDeF Cassandra 0.1.1<sup>74</sup> and Cassandra 1.2.2.<sup>75</sup>

**Ammonium Perchlorate Case Study.** Simulations of AP were performed at 1 atm and 10, 78, and 298 K. Three properties were considered: (1) the absolute percent error (APE) from the experimental lattice parameters averaged across all three temperatures, i.e. the mean absolute percent error (MAPE), and (2) the mean of the absolute residuals of equilibrium average simulated atomic positions in reference to the experimental unit cell<sup>76</sup> at 10 K, subsequently referred to as unit cell mean distance (UCMD), and (3) hydrogen-bonding symmetry that is present in the experimental crystal structure. Four workflow iterations were performed.

**Force Field Parameters.** The Class II force field of Zhu et al.<sup>77</sup> served as a basis for the development of a hand-tuned Class I force field. The partial charges were left unchanged.<sup>78</sup> The Class II intramolecular bonds and angles were recast to the Class I harmonic functional forms; this process was *ad hoc* and involved qualitative matching to the experimental infrared spectrum. The most significant outcome of this procedure was that at 298 K the N–H stretching mode split into two separate peaks for the Class I force field, as opposed to the single peak observed by both experiment and the Class II force field. This is likely due to inherent limitations in the harmonic representation of the vibrational mode; in the context of our work, this trade-off in vibrational behavior for the simplicity and transferability of the Class I AP force field is acceptable. The LJ parameters of the hand-tuned force field were also developed with an *ad hoc* approach, using similar structural metrics as described above. The hand-tuned AP force field parameters are reported in *SI Table S4*.

The force field optimization workflow was applied to further optimize the  $\sigma$  and  $\epsilon$  for the 4 unique atom types in the AP model, giving a total of 8 calibrated parameters. The lower and upper bounds for each parameter were as follows ( $\sigma$  in Å,  $\epsilon$  in kcal/mol):  $3.5 \leq \sigma_{\text{Cl}} \leq 4.5$ ,  $0.5 \leq \sigma_{\text{H}} \leq 2.0$ ,  $2.5 \leq \sigma_{\text{N}} \leq 3.8$ ,  $2.5 \leq \sigma_{\text{O}} \leq 3.8$ ,  $0.1 \leq \epsilon_{\text{Cl}} \leq 0.8$ ,  $0.0 \leq \epsilon_{\text{H}} \leq 0.02$ ,  $0.01 \leq \epsilon_{\text{N}} \leq 0.2$ ,  $0.02 \leq \epsilon_{\text{O}} \leq 0.3$ . The parameter bounds are also summarized in *SI Table S5*. All unlike LJ interactions were calculated with geometric mixing rules.

**Property Calculation Details.** In an effort to be more consistent with the refined hydrogen positions described by Choi et al.,<sup>76</sup> the hydrogen atoms in the primitive cell were extended along their N–H vectors to match the N–H lengths that they report in Table V of their 1974 paper. To assess the symmetry that should be present in orthorhombic AP's *Pnma* space group, the differences in the N–H(3)…O(3) mirror symmetric bond lengths and angles were computed. Hydrogen bonds within 0.001 Å, and angles within 0.3° were considered symmetric. To determine tolerances for assessing symmetry, the manually tuned force field was utilized and the frequency of saving coordinate data over the 100 ps production run was varied between 100 and 10 000 fs. When saving the coordinates every 100 fs, the symmetric hydrogen bond lengths were within 0.000 03 Å and the angles were within 0.01° of each other. When saving the coordinates every 10 000 fs, the resolution of symmetry decreases to within 0.001 Å for bonds and 0.3° for angles. For data management reasons, the coordinates were saved every 10 000 fs and the corresponding symmetry tolerances were utilized in classifying if a given

parameter set was successful in reproducing the experimentally observed symmetry in the hydrogen bonding structure of AP.

**Classifier.** Two SVM classifiers were trained. The first classifier predicted whether a parameter set would yield an accurate 10 K unit cell with  $UCMD < 0.8 \text{ \AA}$ , and the second classifier predicted whether a parameter set would yield the desired hydrogen bond symmetry, as defined above.

**GP Model.** Two GP surrogate models were trained. The first GP model predicted the 10 K  $UCMD$  from  $\zeta'$ . Parameter sets with  $UCMD \geq 0.8 \text{ \AA}$  were not included in the training data. The second GP model predicted the APE of the lattice parameters from  $\zeta'$  and  $T$ . Both GP models used a Matérn  $\nu = 3/2$  kernel and a linear mean function.<sup>66</sup>

**Selecting Parameter Sets for the Next Iteration.** For each iteration, 1 000 000 new parameter sets were generated using LHS. Each parameter set was evaluated with the  $UCMD$  and symmetry classifiers. Parameter sets that did not meet the  $UCMD$  threshold were discarded. The remaining parameter sets were evaluated with the two GP models. The lattice APE GP model was evaluated at  $T = 10, 78$ , and  $298 \text{ K}$  for each parameter set. The mean of the lattice parameter APE at each temperature was calculated and recorded as the lattice MAPE. All parameter sets that did not meet the  $UCMD$  and lattice MAPE thresholds listed in the *SI Table S6* were discarded. When selecting parameter sets for the fourth iteration, the symmetry SVM was used to remove all parameter sets that did not meet the symmetry threshold (*SI Table S6*). A total of 250 parameter sets were selected for the next iteration. All nondominated parameter sets were selected. The remainder of the parameter sets were selected by applying an  $L_1$  distance metric in scaled parameter space and the distance-based search to identify well-separated parameter sets.

**MD Simulations.** Simulations of orthorhombic AP were performed in the  $NpT$  ensemble at 1 atm and 10, 78, and 298 K. The AP structure was taken from the 10 K data of Choi et al.<sup>76</sup> The simulation cell comprised 378 ( $6 \times 9 \times 7$ ) unit cells. Initial velocities were drawn from a Gaussian distribution with the linear and angular momenta set to zero. A 1.0 fs time step was utilized with the time integration scheme derived by Tuckerman et al.<sup>79</sup> The equations of motions were those of Shinoda et al.<sup>80</sup> Nosé–Hoover style algorithms were utilized for both the thermostat and barostat with relaxation times of 0.1 and 1.0 ps, respectively. The  $x$ -,  $y$ -, and  $z$ -dimensions were allowed to fluctuate independently while maintaining an orthorhombic geometry. All simulations utilized 100 ps of equilibration followed by an additional 100 ps for generating production data. Pairwise LJ and Coulombic interactions were computed up to 1.5 nm, and long-range electrostatic interactions were computed using the particle–particle particle–mesh method<sup>70</sup> with a relative accuracy of  $10^{-5}$ . No analytical tail corrections were applied to the repulsion–dispersion interactions. All bonds were fully flexible. Simulations were performed with LAMMPS, version 7 Aug 2019.<sup>81</sup>

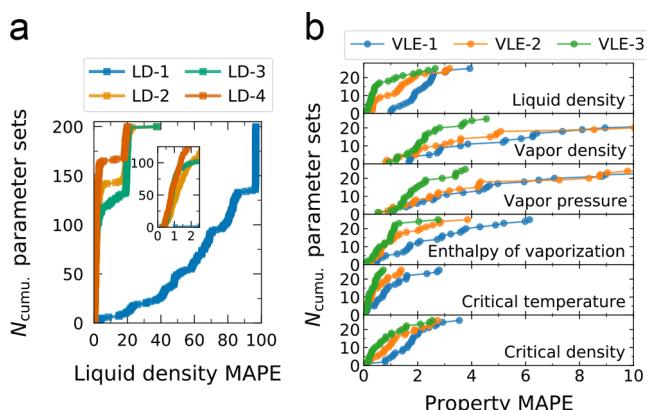
## RESULTS

**Case Study: Hydrofluorocarbon Force Fields.** Recent international agreements, including the 2016 Kigali Amendment to the 1987 Montreal Protocol, mandated the phaseout of high global warming potential HFC refrigerants.<sup>82</sup> Accurate HFC force fields that are compatible with typical all-atom functional forms are of interest as part of a broader multiscale engineering effort to sustainably implement this phaseout.

Here, we optimize force fields for HFC-32 and HFC-125, the two components of R-410a, a common household refrigerant, to accurately predict the pure-component VLE properties. While an accurate hand-tuned force field for HFC-32 exists in the literature,<sup>37</sup> the existing HFC-125 force fields are either inaccurate<sup>15</sup> or rely on less common functional forms,<sup>83–85</sup> which often leads to challenges with force field transferability and simulation software compatibility. For HFC-32, we show that our strategy can develop force fields that outperform expert-created models, while for both HFC-32 and HFC-125, we demonstrate the large improvements that are possible compared against “off-the-shelf” models.

We applied a two-stage approach to improve the HFC force fields. Our workflow was first applied to optimize the force fields to accurately predict the LD at the experimental saturation pressure for five temperatures spanning an 80 K temperature range. Following four iterations (LD-1, LD-2, LD-3, and LD-4), 25 parameter sets with low LD MAPE were used to initiate the second stage of force field optimization. In this stage, force field parameters were optimized to accurately predict VLE properties: saturated liquid density, saturated vapor density, vapor pressure, and enthalpy of vaporization. The two-stage approach has advantages: (1) the MD simulations required to compute LD in the isothermal–isobaric ensemble are computationally less expensive than the MC simulations required to compute VLE properties in the Gibbs ensemble and (2) the stability of the Gibbs ensemble MC simulations is more sensitive to very poor force field parameters.

Figure 2a shows the cumulative number of parameter sets that yield less than some value of the LD MAPE for each HFC-

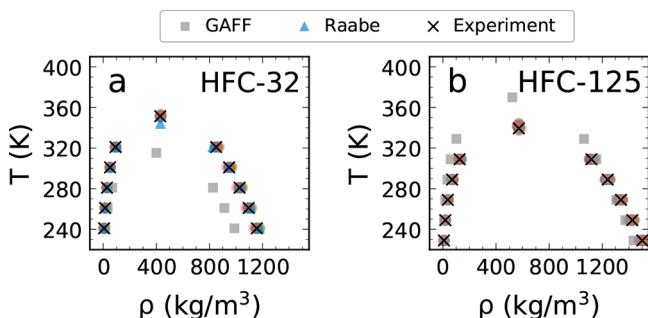


**Figure 2.** Cumulative number of HFC-32 parameter sets generated per iteration with less than some MAPE for (a) the liquid density iterations 1–4 (LD- $n$ ) and (b) vapor–liquid equilibrium iterations 1–3 (VLE- $n$ ), where  $n$  is the iteration number. (a) inset) LD behavior for liquid density MAPE < 2.5%.

32 LD iteration. Analogous results for HFC-125 are reported in *SI Figure S3*. The strength of the surrogate model approach is highlighted by the improvement from the initial liquid density iteration, LD-1, which evaluated 250 parameter sets generated directly from LHS, to the second liquid density iteration, LD-2, which evaluated parameter sets predicted by the surrogate models to yield low LD MAPE. In LD-1 fewer than 5 parameter sets had an LD MAPE below 10%, but LD-2 yielded more than 100 parameter sets with LD MAPE below 2.5%. Limited additional improvements are observed in LD-3 and LD-4, but additional parameter sets with low LD MAPE

are nonetheless generated. Figure 2b shows the same information for three VLE workflow iterations (VLE-1, VLE-2, and VLE-3). Consistent improvements in the saturated liquid density, saturated vapor density, vapor pressure, and enthalpy of vaporization are observed from VLE-1 to VLE-3. The results for the critical temperature and critical density also show improvement even though these properties were not explicitly included in the parameter optimization workflow. Note that the saturated liquid density in VLE-1, which evaluated 25 parameter sets generated during the LD stage, performs slightly worse than the results from LD-4 for two reasons: (1) the model vapor pressure is not precisely equal to the experimental vapor pressure and (2) a smaller system size and shorter interaction cutoff were used to minimize the computational overhead of the LD iterations. Despite the approximation errors introduced by smaller system sizes and cutoffs, the success of our two-stage optimization strategy shows that initial iterations can be performed with less computationally expensive simulations.

After completing the three (five) VLE iterations, our force field parametrization workflow yielded 26 HFC-32 (45 HFC-125) nondominated parameter sets. Figure 3 compares vapor–



**Figure 3.** Vapor–liquid equilibrium envelopes for (a) HFC-32 and (b) HFC-125. The 26 (HFC-32) and 45 (HFC-125) nondominated parameter sets identified in this work are reported as the transparent colored circles and are compared with literature<sup>15,37</sup> and experiment values.<sup>61</sup> All the nondominated parameter sets for both HFCs reproduce the experimental values well and are thus highly overlapped.

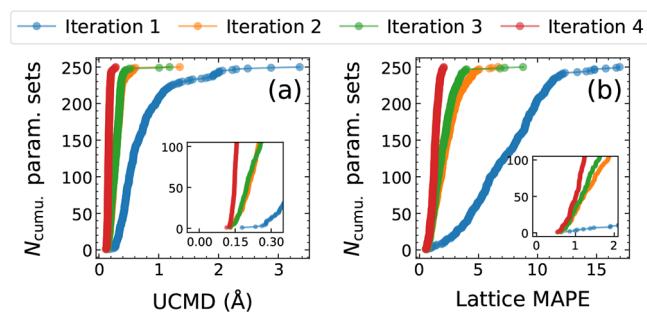
liquid coexistence curves predicted by our nondominated parameter sets with experiments<sup>61</sup> and force fields for HFC-32 and HFC-125 found in the literature. Results for vapor pressure and enthalpy of vaporization are shown in SI Figure S4. The optimized HFC-32 and HFC-125 force fields are notably better than GAFF, and multiple optimized HFC-32 force fields give improved accuracy in all properties compared to the Raabe force field.<sup>37</sup> We chose an error threshold metric to select a subset of top-performing parameter sets from the nondominated sets. This yielded four HFC-32 top parameter sets with MAPE of less than 1.5% and four HFC-125 top parameter sets with MAPE of less than 2.5% for the four properties included in the optimization workflow and the critical temperature and critical density. Comparisons of critical temperature and critical density values between experiment, the top four optimized force fields, and literature force fields for both HFCs are shown in SI Tables S7 and S8.

**Case Study: Ammonium Perchlorate Force Field.** AP is a key ingredient in some solid rocket propellants. Experimental data for physical properties of AP are readily available and a Class II force field parametrized by Zhu et al.<sup>77</sup>

has been used to predict<sup>78</sup> pure AP properties at temperatures up to 298 K. The Class II functional form supplements the harmonic diagonal constants found in the more common Class I force fields through the inclusion of cross terms, namely, the stretch–stretch and stretch–bend interactions. The cross terms couple internal coordinates in an effort to better reproduce the molecular energetics as well as the dynamics of a system by accounting for anharmonic and coupling interactions. However, it is of interest to develop a Class I force field for AP to use in conjunction with existing Class I force fields for the other components of conventional solid propellant, aluminum oxide<sup>86</sup> and the polymeric binder.<sup>87</sup> Here, we parametrize an AP force field with our force field optimization workflow; we previously had utilized hand-tuning methods to develop a Class I AP force field. We present a comparison between the conventional hand-tuning approach and our workflow. In addition to the motivation provided above, we selected solid AP as our second case study because it represents a very different system than the HFC VLE investigated in the first case study.

The properties to which we calibrated our Class I force field were the following: (1) UCMD at 10 K, defined as the mean of the absolute residuals of equilibrium average simulated atomic positions in reference to the experimentally observed unit cell atomic positions (low values indicate the simulation maintains the experimental AP crystal structure); (2) unit cell lattice parameter mean absolute percent error at the three temperatures of interest (10, 78, and 298 K); and (3) correct hydrogen bond symmetry.

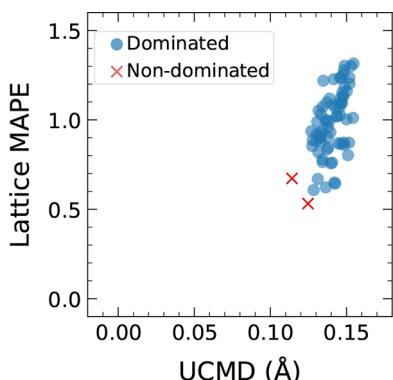
Four iterations of the force field optimization workflow were performed. The cumulative error plots are shown in Figure 4.



**Figure 4.** Cumulative number of AP parameter sets per iteration with less than some value of (a) the 10 K unit cell mean distance (UCMD) and (b) the lattice MAPE. Insets have the same axis titles and focus on the improvement from iteration 3 to iteration 4. Less strict UCMD and lattice MAPE criteria were applied when selecting parameter sets for iterations 2 and 3, and stricter criteria were applied when selecting parameter sets for iteration 4. Threshold values for selecting next iteration points are shown in SI Table S6.

Once again, we observe substantial improvement between the first and second workflow iteration. Here, the cumulative error plots also show that the criteria for selecting parameter sets for the next iteration can significantly affect the improvement in objective performance between iterations. Less strict UCMD and lattice MAPE criteria were applied when selecting parameter sets for iterations 2 and 3, and stricter criteria were applied when selecting parameter sets for iteration 4; iteration 4 showed much greater improvement over iteration 3 whereas iterations 2 and 3 are very similar. Our workflow generated 70 parameter sets over the four iterations which gave

lower UCMD and lattice parameter errors than the hand-tuned values while maintaining the correct hydrogen bonding symmetry. We found two nondominated parameter sets, as shown in Figure 5. These two nondominated parameter sets



**Figure 5.** 70 AP parameter sets that yield lower UCMD and lattice parameter errors than the hand-tuned values while maintaining the correct hydrogen bonding symmetry. The red points are nondominated and indicate our top two AP parameter sets. The blue points are dominated.

will subsequently be referred to as our top two AP parameter sets. Table 1 compares the AP results for these top parameter sets with the hand-tuned and Class II force field results.

**Table 1. Crystal Structure Results for the Top Two AP Parameter Sets, “Top A and Top B”, Identified via the Workflow Presented in This Study, the Hand-Tuned Parameter Set (HT), and the Class II Parameter Set of Zhu et al.<sup>77a</sup>**

| property      | T (K) | Top A  | Top B  | HT     | Class II |
|---------------|-------|--------|--------|--------|----------|
| lat. <i>a</i> | 298   | -0.77  | -0.40  | -2.09  | -0.21    |
|               | 78    | -0.88  | -0.48  | -1.87  | -2.79    |
|               | 10    | -0.24  | 0.26   | -1.38  | -3.10    |
| lat. <i>b</i> | 298   | 1.13   | 0.61   | 1.96   | 7.00     |
|               | 78    | 1.11   | 0.89   | 1.68   | 8.19     |
|               | 10    | 0.63   | 0.26   | 1.16   | 8.22     |
| lat. <i>c</i> | 298   | -0.18  | -0.74  | -1.04  | 1.64     |
|               | 78    | -0.71  | -1.10  | -1.31  | 0.46     |
|               | 10    | 0.39   | 0.04   | -0.30  | 0.32     |
| MAPE          |       | 0.67   | 0.53   | 1.42   | 3.55     |
| UCMD          | 10    | 0.1142 | 0.1247 | 0.1560 | 0.3485   |

<sup>a</sup>Lattice parameter results are reported in terms of percent error relative to experimental results.<sup>76</sup> The UCMD results are given in angstroms.

## ■ DISCUSSION

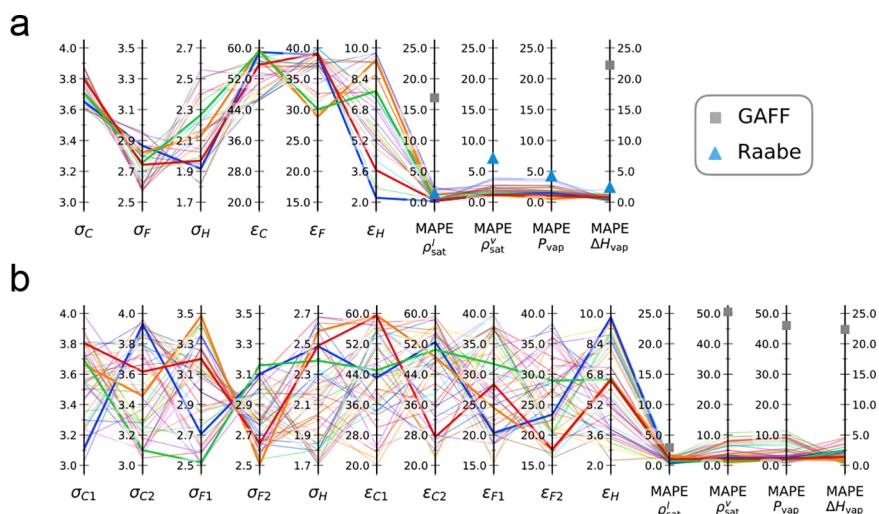
**Many Distinct Parameter Sets Yield Equally Accurate Results.** The conventional wisdom in molecular modeling often seems to be that there is a single “correct” or “best” set of force field parameters, but this may be a misleading way to think about force field optimization. No force field is a perfect representation of the physical world. Therefore, model limitations will result in trade-offs between different objectives and, depending on the property priorities for a specific application, lead to different optimal parameter sets.<sup>88</sup> However, our results clearly show that multiple parameter sets can reproduce several experimental properties with very

low error. For the HFCs, our procedure yielded 26 (HFC-32) and 45 (HFC-125) nondominated parameter sets, which are distinctly different parametrizations, all of which display good performance on our optimization objectives and the critical temperature and density. A visual representation of the nondominated parameter sets and their performance for the optimization objectives is shown in Figure 6. For HFC-32, where there are six optimized force field parameters, the nondominated parameter sets show variation of up to  $\sim 0.3$  Å in the carbon and fluorine  $\sigma$  values and up to  $\sim 10$  K/ $k_B$  in the carbon and fluorine  $\epsilon$  values. For HFC-125, there is even larger variation in the  $\sigma$  and  $\epsilon$  values among the nondominated parameter sets. We suspect this is because there are a larger number of parameters for HFC-125 (10) than for HFC-32 (6), allowing for compensating behavior between different parameters. For example, consider  $\sigma_{F1}$  and  $\sigma_{F2}$ . There is a clear compensating effect: when  $\sigma_{F1}$  is larger,  $\sigma_{F2}$  is smaller, and vice versa. On the other hand,  $\sigma_{F1}$  and  $\sigma_{F2}$  do appear to be different, as some parametrizations of  $\sigma_{F1}$  are 0.3 Å larger than any of the parametrizations of  $\sigma_{F2}$ .

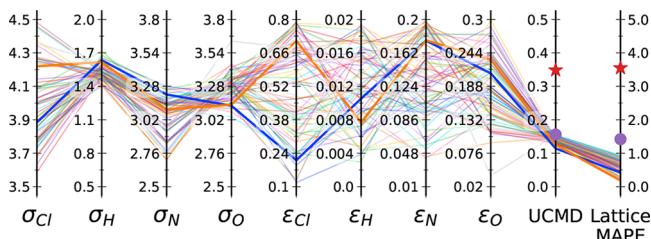
The visualizations in Figure 6 suggest that the 26 (HFC-32) and 45 (HFC-125) nondominated parameter sets are indeed distinct parametrizations, rather than closely related parametrizations with small variations along a continuous manifold of good parameters. To further investigate this question, the  $L_1$  distance between the best-performing parameter set in each property and every other nondominated parameter set was calculated and plotted against the property error (SI Figure S5). No correlation is observed between the similarity of a parameter set to the top-performing parameter set in a given property and the property error for that parameter set. This strongly suggests that our nondominated parameter sets are indeed distinct parametrizations. In part, this can be attributed to our procedure for advancing parameter sets to the next iteration, where we intentionally selected points that were well-separated in parameter space (see section *Selecting Parameter Sets for the Next Iteration in Hydrofluorocarbon Case Study*).

Similar behavior is observed in the AP system, where we identified 70 parameter sets that outperform the hand-tuned Class I and existing Class II force fields.<sup>77</sup> Figure 7 shows the variation in the optimized AP force field parameters. Once again, a number of distinct parametrizations yield similar accuracy for the optimization objectives. The  $\sigma$  values vary by  $\sim 0.3$  Å for the hydrogen and oxygen atom types that are more exposed to intermolecular interactions, and up to as much as nearly 1.0 Å for the buried Cl atom type. The  $\epsilon$  values vary by as much as  $\sim 0.6$  kcal/mol, with the largest variation once again observed for the Cl atom type. Although there is a large variation in the individual parameter values between different parameter sets, it is the entire parameter set, taken together, that provides good performance. The results presented here do not suggest that a parameter can take any value within the ranges shown in Figure 7, e.g., any value of  $\sigma_{Cl}$  between 3.5 and 4.5 Å, and yield good performance if all other other parameter values are held constant. Rather, correlations between the different parameters enable a number of distinct yet highly accurate force field parametrizations.

Finding many distinct well-performing nondominated parameter sets suggests the model may be overparameterized. To investigate this, we performed a local identifiability analysis by inspecting the eigenvalues of the Fisher information matrix (FIM) for the top four parameter sets for both the HFC-32 and HFC-125 models. As detailed in the SI Discussion, we find



**Figure 6.** Repulsion–dispersion parameters for (a) 26 HFC-32 and (b) 45 HFC-125 high quality parameter sets.  $\sigma$  is reported in units of Å, and  $\epsilon$  is reported in units of  $K/k_B$ . Each parameter set is connected by a different color line. Thick lines indicate the top four parameter sets for each molecule. The y-axes are scaled to show the full range investigated for each parameter. The final four y-axes show the performance for the training objectives. The gray squares and cyan triangles show the performance of GAFF<sup>15</sup> and the force field of Raabe,<sup>37</sup> respectively. For HFC-32 the GAFF MAPE for  $\rho_{\text{vap}}$  and  $P_{\text{vap}}$  are not shown as they are 133 and 104, respectively.

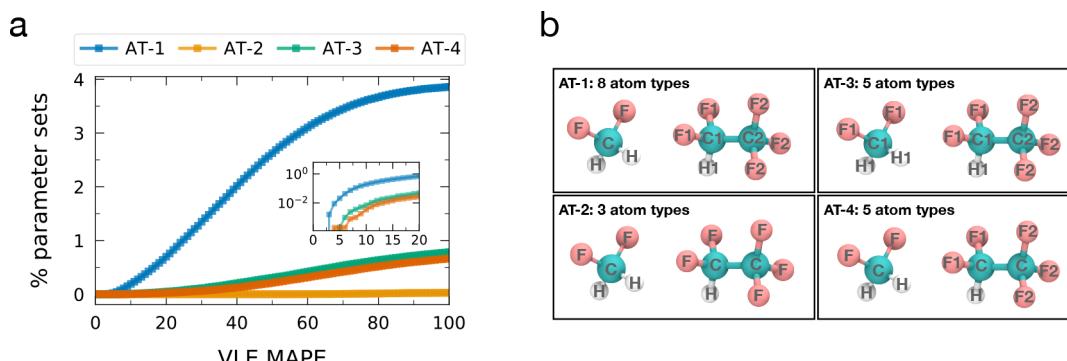


**Figure 7.** Repulsion–dispersion parameters for the final 70 AP parameter sets.  $\sigma$  is reported in units of Å, and  $\epsilon$  is reported in units of kcal/mol. Each parameter set is connected by a different color line. The thick lines show the top two AP parameter sets. The y-axes are scaled to show the full range investigated for each parameter. The final two y-axes show the training objectives. The red stars and purple circles show the performance of the Class II force field of Zhu et al.<sup>77</sup> and the hand-tuned Class I force field, respectively.

the FIM has one and five near-zero eigenvalues for HFC-32 and HFC-125, respectively, when considering only the liquid density data. This means we can only identify five (HFC-32:6 total parameters minus 1 near-zero eigenvalue equals five

identifiable directions, HFC-125:10 minus 5 equals 5) parameters using only experimental liquid density data. The corresponding eigenvectors for these near zero eigenvalues reveal the direction in parameter space in which the regression objective is flat (near zero curvature). Unfortunately, these eigenvectors do not point in the direction of a single parameter, which complicates their interpretation. More importantly, the FIM is full rank when simultaneously regressing both liquid density and VLE experimental data sets, which implies both models are locally fully identifiable. Thus, this analysis resolves one aspect of overparameterization by mathematically quantifying the importance of including multiple types of experimental data in the model calibration process. Moreover, our results suggest all of the top parameter sets are near locally optimal solutions (all with positive curvature, thus locally identifiable).

Another aspect of overparameterization is that we find a large number of high-quality solutions. These results are not surprising, given that many inverse problems based on engineering models have numerous locally optimal parameter sets that lead to accurate in-sample predictions.<sup>89</sup> In this case,



**Figure 8.** (a) Cumulative percent of parameter sets from a large ( $O(10^5)$ ) Latin hypercube that yield less than each value of MAPE for all four VLE properties. For a given MAPE, a higher percentage indicates that more parameterizations achieve at least that threshold level of accuracy. Results are shown for four different atom-typing schemes (AT-1, AT-2, AT-3, and AT-4). (inset) Low-MAPE region with data on a log-scale. (b) Schematic of AT-1, AT-2, AT-3, and AT-4. AT-1 is the original atom-typing scheme where no atom types were shared between HFC-32 and HFC-125.

we hypothesize that parametrizing each molecule individually leads to many locally optimal parameter sets. Extending our method to simultaneously optimize force field parameters for an entire class of molecules (e.g., all hydrofluorocarbons) with a number of shared atom types will likely reduce the overparameterization. While we leave the development of an HFC force field for future work, here, we explore the effects of using shared atom types for HFC-32 and HFC-125 on the number of high-quality model parametrizations. We consider four atom-typing schemes (AT-1, AT-2, AT-3, and AT-4), shown in Figure 8b. AT-1 is the scheme we have used thus far; there are eight total atom types, three for HFC-32 and five for HFC-125. In AT-2, we use a total of three atom types across both molecules, C, F, and H. AT-3 and AT-4 both use five atom types but differ in how these atom types are distributed. In AT-3, we maintain the original scheme for HFC-125, but then reuse the C1, F1, and H1 types for HFC-32. In AT-4, the C and H types are shared as they are either small or buried, while each fluorine is a different atom type. The surrogate models trained during this work were used to evaluate the performance of the different atom typing schemes. LHS was used to generate 500 000 parameter sets. First, the liquid density GP surrogate model was used to eliminate any parameter sets with RMSE greater than 100 kg/m<sup>3</sup>. For each of the remaining parameter sets, the VLE GP surrogate models were used to predict the MAPE for each VLE property (saturated liquid and vapor densities, vapor pressure, and enthalpy of vaporization). Figure 8a reports the percentage of the original 500 000 parameter sets that yield less than a given MAPE threshold for all four VLE properties, simultaneously. The atom-typing schemes with a reduced number of atom types have a much smaller percentage of parameter space containing low-error parameter sets. In fact, AT-2, with only three atom types, does not result in any parametrizations that are predicted to have below 46% MAPE for all four VLE properties. AT-3 and AT-4 show that even with the same number of atom types, one atom-typing scheme may result in superior performance. This naturally raises another question: given different atom-typing schemes, which should be used? Recent work<sup>90</sup> demonstrates the promise of using Bayes factors to compare models with different levels of complexity (e.g., different atom-typing schemes) and make a justified selection.

Since the prior analysis was performed entirely with the predictions of the GP surrogate models, we performed molecular simulations with two top-performing parameter sets for each of the shared atom-typing schemes (AT-2, AT-3, and AT-4) in order to compute the simulated MAPE values and compare them with the surrogate model predictions. The results are reported in SI Table S9. Overall, the surrogate model predictions were excellent, often showing less than 0.5% MAPE deviation from the simulated MAPE. GEMC simulations for AT-2 were unstable at the highest temperature, confirming the surrogate models' prediction that AT-2 would not yield any good parameter sets. We also explored HFC-125-only force fields with a reduced number of atom types (SI Table S10) and found that we were able to identify parameter sets with less than 3% MAPE using only 3 atom types (C, F, and H). However, as noted above, when we attempted to use three atom types (C, F, and H) for both HFC-32 and HFC-125, no good force fields were identified. This finding is strong evidence that the fluorine atom types in HFC-32 and HFC-125 should be different (e.g., AT-4) and shows how developing

parametrizations for an entire class of molecules will reduce the number of viable parameter sets.

Adding additional objective properties is a complementary strategy to reduce the number of viable parameter sets. In that case, it is important that the additional properties are orthogonal in the sense that good performance for one property is not highly correlated with good performance for another property. If property performance is highly correlated, then adding additional properties to the optimization workflow may not substantially reduce the number of viable parameter sets. The apparent overparameterization observed in this work emphasizes why tuning force fields for specific systems and using a few objective properties via relatively simple methods such as epsilon-scaling, manipulating mixing rules, or varying a single parameter value are often quite successful. However, our findings suggest that the force fields developed via these methods are most likely only one of a large number of possible parametrizations that would yield at least equal accuracy.

A further question involves how final parameter sets should ultimately be selected, given that many high-quality parameter sets are available. Our workflow is explicitly not designed to identify a single optimal set of force field parameters. Instead, it searches for and identifies high quality parameter sets with respect to all of the optimization objectives, e.g., points in the nondominated set. Selecting a single specific parameter set from the optimized parameter sets identified by the workflow requires additional *post hoc* criteria that are application specific. Here, we chose nondominated status and error thresholds for all properties. Alternative strategies include creating a weighted sum of errors in the properties based upon the desired application and domain knowledge, ranking force fields by their error in the various properties studied via statistical tests,<sup>54</sup> evaluating the force field's performance for properties not included in the optimization procedure, or selecting parameter sets based upon a measure of compatibility with the force fields being used for other components of a system. One could also consider chemical intuition when selecting the final parameter sets, e.g., for HFC-125, perhaps a parameter set with more similar values for both fluorine atoms would be preferred. Though our preference is to minimize the number of *ad hoc* choices, ultimately, selecting the final force field for a given application will be system and application dependent and rely heavily on domain expertise.

**Maintaining a Physically-Motivated Analytical Functional Form Aids Transferability to Properties Not Included as Optimization Objectives.** One important question is whether the force field parameters developed with this workflow will yield accurate property predictions for properties not included in the optimization workflow. We have already shown that the HFC force fields developed during the VLE tuning stage result in accurate critical temperature and density even though these properties were not optimization objectives. However, these critical properties are largely determined by accurately capturing the temperature dependence of the saturated liquid and saturated vapor density, both of which were optimization objectives. To further investigate the transferability of force field parameters developed with our workflow to properties not included as optimization objectives, we examine the performance of the 25 parameter sets used during the VLE-1 iteration. These parameter sets were used for VLE-1 because they were identified as good at predicting the temperature dependence of the liquid density during the LD iterations. Figure 2 shows that when applied for VLE-1, many

perform quite well for VLE properties. In fact, three of the HFC-32 parameter sets used for the VLE-1 iteration had less than 2% MAPE in all six properties. Furthermore, when compared with GAFF, all 25 parameter sets selected from the LD stage yield better performance for all six properties. This is strong evidence that our force field optimization workflow can, with the correct optimization objectives, yield force fields that accurately predict properties beyond the optimization objectives.

The transferability of the LD-optimized parameters to VLE gives credence to our overall force field optimization philosophy, which maintains traditional analytical functional forms and uses machine learning as a guide to identify optimal parameters. However, *a priori*, it is unclear that there should be such a strong correlation between the liquid density and VLE properties. For many systems, accurately predicting the liquid density is a necessary, but often quite insufficient, condition for an accurate force field. We hypothesize there is a key factor that contributes to the transferability of the parameters developed during the LD iterations to VLE: the LD simulations were performed at the saturated vapor pressure across an 80 K temperature range, up to within 30 K of the experimental critical temperature. Accurately capturing the liquid density at saturation across a relatively large temperature range and avoiding spontaneous vaporization, especially at conditions closer to the critical point, requires capturing a careful balance of the cohesive energy and molecular size, which are closely related to the LJ repulsion–dispersion parameters that were calibrated. If the correlation between LD-optimized parameters and VLE properties proves applicable to other classes of molecules, it may offer a rapid method for developing force fields with accurate VLE properties.

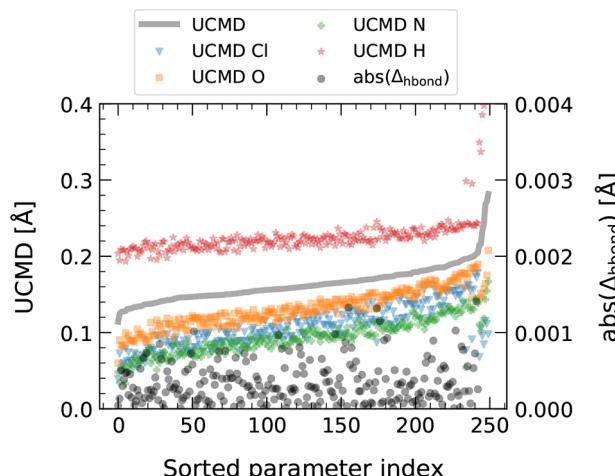
**Selecting Good Properties for Force Field Optimization is Challenging.** When optimizing force fields for the HFC case study, we were interested in developing force fields that accurately predict HFC VLE behavior. As such, we chose to calibrate parameters to the saturated liquid and vapor densities, vapor pressure, and enthalpy of vaporization. However, these properties are expensive to compute in molecular simulations, making it difficult to evaluate a large parameter space. Therefore we used less computationally expensive LD iterations to generate good parameter sets for VLE and narrow the parameter search space. Furthermore, we continued to use the highly accurate LD GP surrogate models to screen out poor parameter sets during the VLE iterations. The success of this approach demonstrates that a cheaper “screening” property can be used to narrow the parameter search space drastically when good parameter sets for the screening property are a superset of the good parameter sets for the final properties of interest.

The AP case study had different challenges. The MD simulations required to predict the AP properties were computationally inexpensive, so there was no need to first use a screening property. However, it was not immediately clear what experimental properties we should target. Our first implementation attempted to reproduce the temperature dependence of the crystal lattice parameters alone; this proved ineffective and naive in hindsight, as we generated many force fields that yielded the correct crystal lattice parameters but incorrect crystal structures. To overcome this issue, we added the 10 K UCMD as an objective because it is a measure of how accurately the force field reproduces the experimental crystal structure at 10 K. The lattice MAPE was still included to

capture the temperature dependence of the crystal dimensions since the experimental unit cell coordinates are only reported at 10 K.

The UCMD surrogate model has a notable difference from the others; whereas the other surrogate models predict a property (e.g., lattice  $a$  or  $p_{\text{vap}}$ ), the UCMD is itself an objective function. The UCMD surrogate model predicts the mean distance of all of the unit cell atoms from their respective coordinates in the experimental unit cell. By definition, this distance is zero if the simulated structure perfectly matches experiment. There are benefits to using physical experimentally measured properties compared to an objective function within the optimization workflow, including providing a clear mapping between a surrogate model and the objective metric. However, using surrogate models to predict the value of an objective function provides the opportunity to combine multiple pieces of information into a single quantity, as is the case with UCMD, which combines the distance of 40 atoms from their positions in the experimental unit cell into a single value. This strategy can drastically reduce the number of required surrogate models. In general, our experience with the AP case study emphasizes that careful thought must be given as to which experimental properties are best to target and how these should be accounted for within the workflow. Roughly 75% of our effort for the AP case study was dedicated to identifying the appropriate experimental properties to target.

**Systematic Parameter Search Provides Insights into Model Limitations.** The exhaustive search of parameter space enabled by our workflow provides opportunities to distinguish between inaccurate results from poor parameter sets and physical limits from our choice in force field functional form and unoptimized parameters. For example, although our workflow finds high-quality AP parameter sets, we encountered limitations that likely arise from parameters that were not calibrated, and possibly even the force field functional form that we selected. No parameter set predicted an overall UCMD of less than 0.1 Å. Given the exhaustive search enabled by our force field optimization workflow, this suggests that there are no parameter sets capable of yielding a crystal structure with UCMD below 0.1 Å, given the selected functional form, intramolecular parameters, and partial charges. Figure 9 shows the per-element UCMD distances after iteration 4. Although the UCMD for the chlorine, oxygen, and nitrogen atoms fall below 0.12 Å for many parameter sets, the hydrogen UCMD rarely falls below 0.2 Å. Further investigation suggests that this effect is because the N–H bond stretching is insufficiently susceptible to the three unique local hydrogen-bonding chemical environments; experiments report<sup>76</sup> that the N–H bond lengths range between 1.028 and 1.058 Å whereas in simulations the N–H bond lengths typically cover a much smaller range—between 1.025 and 1.033 Å—for parameter sets that well reproduce the experimental physical properties. The N–H stretching force constant was not included in our parametrization process. However, even if it was, it is not clear that it would be possible to capture the correct bond stretching behavior and match the vibrational spectra and the N–H bond lengths with a Class I functional form. The exhaustive search provides confidence that the limitations of the model arise from the functional form and unoptimized parameters, rather than the selected parametrization.



**Figure 9.** Overall (gray line) unit cell mean distance (UCMD) compared with UCMD of the four different atom types (points) for the parameter sets tested during iteration four of the ammonium perchlorate force field optimization. The hydrogen bond symmetry is reported as  $\text{abs}(\Delta_{\text{hbond}})$ , where  $\Delta_{\text{hbond}}$  is the difference in the symmetric hydrogen bond lengths.

## CONCLUSIONS

We have presented a machine learning directed workflow for top-down optimization of force field parameters. By harnessing surrogate-assisted optimization, our workflow drastically reduces the number of simulations necessary to find optimal force field parameters by replacing them with computationally tractable surrogate model evaluations. We synthesize GPR and SVM surrogate models and multiobjective optimization into a generic approach to optimize all-atom force fields for realistic systems. We have applied our workflow to optimize HFC force fields to VLE properties and an AP force field to the experimental crystal structure. These case studies show that our workflow can be used for systematic exhaustive screening of parameter space and that surrogate models are highly effective at predicting both simulated physical properties and objective metrics, enabling us to find multiple low-error force fields. The approach presented here could be further combined with gradient-based methods or other approaches such as trust region surrogate-based optimization<sup>91</sup> to further refine the final force fields.

Based upon the success of our approach for the two disparate case studies presented here, we believe that this workflow can be applied to most molecular systems and optimization objectives, provided sufficient reference data. Surrogate models could be used to predict difficult-to-compute thermodynamic properties such as solubilities and binding energies, and transport properties such as self-diffusivity and thermal conductivity. While we have focused on calibrating repulsion–dispersion parameters in this work, this workflow could be used to calibrate any parameters within the force field in a fully top-down approach or as part of a bottom-up force field development workflow, by including *ab initio* data in the fitting procedure.<sup>45</sup> Additionally, we discussed the reasons for successes and limitations of the workflow, the potential challenges of applying this workflow to a particular system (i.e., choosing optimization objectives) and the questions about molecular modeling these results present. We highlight that this workflow is built on a foundation of domain knowledge in selecting the parameters to calibrate, the

parameter bounds, and the experimental properties to ensure results are reasonable.

Finally, while we believe that our workflow will enable more efficient force field development and optimization in the future, reducing the need for laborious hand-tuning practices, quantifying the workflow's efficiency was beyond the scope of this work. We can, however, anecdotally note for the AP case study that the hand-tuning approach utilized  $\sim 15\,000$  simulations and only found 1 optimal parameter set. This is in contrast to our presented workflow, which evaluated  $\sim 3\,000\,000$  parameter sets using surrogate models,  $\mathcal{O}(10^3)$  times as many as the hand-tuning method, but only required 3000 simulations, to find 70 parameter sets with lower error in the metrics of interest than the hand-tuned parameter set. We anticipate that further refining the proposed workflow, e.g., incorporating adaptive sampling via Bayesian optimization, can dramatically reduce the number of molecular simulations required to identify parameter sets that accurately predict several physical properties.

## DATA AND SOFTWARE AVAILABILITY

Codes used to perform the HFC case study and all generated parameter sets are available at <https://github.com/dowlinglab/hfcs-ffit>. Codes used to perform the AP case study and all generated parameter sets are available at <https://github.com/dowlinglab/ap-ffit>.

## ASSOCIATED CONTENT

### Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jcim.1c00448>.

Figures demonstrating the accuracy of the GP models, parameter improvement for HFC-125 during the LD and VLE iterations, performance of the optimized force fields for vapor pressure and enthalpy of vaporization, and correlation between the distance from the top parameter set for each property and the property error. Tables reporting values for all force field parameters not included in the optimization workflow, parameter bounds for optimized parameters, screening criteria for AP iterations, comparisons of the predicted HFC critical temperatures and densities with experiment, and performance of force fields with shared atom types (PDF)

Spreadsheets with eigenvalues and eigenvectors from the identifiability analysis for HFC-32 (ZIP)

Spreadsheets with eigenvalues and eigenvectors from the identifiability analysis for HFC-125 (ZIP)

## AUTHOR INFORMATION

### Corresponding Author

Edward J. Maginn – Department of Chemical and Biomolecular Engineering, University of Notre Dame, Notre Dame, Indiana 46556, United States;  [orcid.org/0000-0002-6309-1347](https://orcid.org/0000-0002-6309-1347); Email: [ed@nd.edu](mailto:ed@nd.edu)

### Authors

Bridgette J. Befort – Department of Chemical and Biomolecular Engineering, University of Notre Dame, Notre Dame, Indiana 46556, United States;  [orcid.org/0000-0002-2340-6092](https://orcid.org/0000-0002-2340-6092)

Ryan S. DeFever – Department of Chemical and Biomolecular Engineering, University of Notre Dame, Notre Dame, Indiana 46556, United States;  [orcid.org/0000-0001-5311-6718](https://orcid.org/0000-0001-5311-6718)  
Garrett M. Tow – Department of Chemical and Biomolecular Engineering, University of Notre Dame, Notre Dame, Indiana 46556, United States;  [orcid.org/0000-0001-5191-6092](https://orcid.org/0000-0001-5191-6092)  
Alexander W. Dowling – Department of Chemical and Biomolecular Engineering, University of Notre Dame, Notre Dame, Indiana 46556, United States;  [orcid.org/0000-0001-9383-7499](https://orcid.org/0000-0001-9383-7499)

Complete contact information is available at:  
<https://pubs.acs.org/10.1021/acs.jcim.1c00448>

## Author Contributions

<sup>‡</sup>B.J.B. and R.S.D. contributed equally to this work.

## Notes

The authors declare no competing financial interest.

## ACKNOWLEDGMENTS

R.S.D. and E.J.M. acknowledge funding from the National Science Foundation, Award number OAC-1835630. B.J.B., A.W.D., and E.J.M. acknowledge funding from the National Science Foundation, Award number CBET-1917474. B.J.B. acknowledges funding from the Richard and Peggy Notebaert Premier Fellowship. G.M.T. and E.J.M. acknowledge funding from the Air Force Office of Scientific Research under Contract AFOSR FA9550-18-1-0321. Computational resources were provided by the Notre Dame Center for Research Computing.

## ABBREVIATIONS

AP, ammonium perchlorate; APE, absolute percent error; AT, atom-typing scheme; FIM, Fischer information matrix; GAFF, general AMBER force field; GEMC, Gibbs ensemble Monte Carlo; GP, Gaussian process; GPR, Gaussian process regression; HFC, hydrofluorocarbon; HFC-125, pentafluoroethane; HFC-32, difluoromethane; LD, liquid density; LHS, Latin hypercube sampling; LJ, Lennard-Jones; MAPE, mean absolute percent error; MC, Monte Carlo; MD, molecular dynamics; RMSE, root mean square error; SVM, support vector machine; UCMD, unit cell mean distance; VLE, vapor-liquid equilibrium

## REFERENCES

- (1) Maginn, E. J. From Discovery to Data: What Must Happen for Molecular Simulation to Become a Mainstream Chemical Engineering Tool. *AIChE J.* **2009**, *55*, 1304–1310.
- (2) Iftimie, R.; Minary, P.; Tuckerman, M. E. *Ab initio* Molecular Dynamics: Concepts, Recent Developments, and Future Trends. *Proc. Natl. Acad. Sci. U. S. A.* **2005**, *102*, 6654–6659.
- (3) Heo, L.; Feig, M. Experimental Accuracy in Protein Structure Refinement via Molecular Dynamics Simulations. *Proc. Natl. Acad. Sci. U. S. A.* **2018**, *115*, 13276–13281.
- (4) Hollingsworth, S. A.; Dror, R. O. Molecular Dynamics Simulation for All. *Neuron* **2018**, *99*, 1129–1143.
- (5) Franco, A. A.; Rucci, A.; Brandell, D.; Frayret, C.; Gaberscek, M.; Jankowski, P.; Johansson, P. Boosting Rechargeable Batteries R&D by Multiscale Modeling: Myth or Reality? *Chem. Rev.* **2019**, *119*, 4569–4627.
- (6) Youngs, T. G. A.; Del Pópolo, M. G.; Kohanoff, J. Development of Complex Classical Force Fields through Force Matching to *ab initio* Data: Application to a Room-Temperature Ionic Liquid. *J. Phys. Chem. B* **2006**, *110*, 5697–5707.

- (7) Lobanova, O.; Mejía, A.; Jackson, G.; Müller, E. A. SAFT- $\gamma$  Force Field for the Simulation of Molecular Fluids 6: Binary and Ternary Mixtures Comprising Water, Carbon Dioxide, and *n*-Alkanes. *J. Chem. Thermodyn.* **2016**, *93*, 320–336.
- (8) Noé, F.; Tkatchenko, A.; Müller, K.-R.; Clementi, C. Machine Learning for Molecular Simulation. *Annu. Rev. Phys. Chem.* **2020**, *71*, 361–390.
- (9) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chem. Rev.* **2021**, DOI: [10.1021/acs.chemrev.0c01111](https://doi.org/10.1021/acs.chemrev.0c01111).
- (10) Behler, J.; Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.* **2007**, *98*, 146401.
- (11) Schütt, K. T.; Arbabzadah, F.; Chmiela, S.; Müller, K. R.; Tkatchenko, A. Quantum-Chemical Insights from Deep Tensor Neural Networks. *Nat. Commun.* **2017**, *8*, 1–8.
- (12) Harrison, J. A.; Schall, J. D.; Maskey, S.; Mikulski, P. T.; Knippenberg, M. T.; Morrow, B. H. Review of Force Fields and Intermolecular Potentials Used in Atomistic Computational Materials Research. *Appl. Phys. Rev.* **2018**, *5*, 031104.
- (13) Rappe, A. K.; Casewit, C. J.; Colwell, K. S.; Goddard, W. A.; Skiff, W. M. UFF, A Full Periodic Table Force Field for Molecular Mechanics and Molecular Dynamics Simulations. *J. Am. Chem. Soc.* **1992**, *114*, 10024–10035.
- (14) 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.
- (15) 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.
- (16) Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, I.; Mackerell, A. D., Jr. CHARMM General Force Field: A Force Field for Drug-Like Molecules Compatible with the CHARMM All-Atom Additive Biological Force Fields. *J. Comput. Chem.* **2009**, *31*, 671–690.
- (17) Martin, M. G. Comparison of the AMBER, CHARMM, COMPASS, GROMOS, OPLS, TraPPE and UFF Force Fields for Prediction of Vapor-Liquid Coexistence Curves and Liquid Densities. *Fluid Phase Equilib.* **2006**, *248*, 50–55.
- (18) Bayly, C. I.; Cieplak, P.; Cornell, W.; Kollman, P. A. A Well-Behaved Electrostatic Potential Based Method Using Charge Restraints for Deriving Atomic Charges: The RESP Model. *J. Phys. Chem.* **1993**, *97*, 10269–10280.
- (19) Wang, J.; Kollman, P. A. Automatic Parameterization of Force Field by Systematic Search and Genetic Algorithms. *J. Comput. Chem.* **2001**, *22*, 1219–1228.
- (20) Barone, V.; Cacelli, I.; De Mitri, N.; Licari, D.; Monti, S.; Prampolini, G. Joyce and Ulysses: Integrated and User-Friendly Tools for the Parameterization of Intramolecular Force Fields From Quantum Mechanical Data. *Phys. Chem. Chem. Phys.* **2013**, *15*, 3736–3751.
- (21) Betz, R. M.; Walker, R. C. Paramfit: Automated Optimization of Force Field Parameters for Molecular Dynamics Simulations. *J. Comput. Chem.* **2015**, *36*, 79–87.
- (22) Zahariev, F.; De Silva, N.; Gordon, M. S.; Windus, T. L.; Dick-Perez, M. ParFit: A Python-Based Object-Oriented Program for Fitting Molecular Mechanics Parameters to *ab initio* Data. *J. Chem. Inf. Model.* **2017**, *57*, 391–396.
- (23) Wang, L.-P.; McKiernan, K. A.; Gomes, J.; Beauchamp, K. A.; Head-Gordon, T.; Rice, J. E.; Swope, W. C.; Martínez, 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.
- (24) Ivanov, M. V.; Talipov, M. R.; Timerghazin, Q. K. Genetic Algorithm Optimization of Point Charges in Force Field Develop-

ment: Challenges and Insights. *J. Phys. Chem. A* **2015**, *119*, 1422–1434.

(25) Wang, R.; Ozhgibesov, M.; Hirao, H. Analytical Hessian Fitting Schemes for Efficient Determination of Force-Constant Parameters in Molecular Mechanics. *J. Comput. Chem.* **2018**, *39*, 307–318.

(26) Grimme, S. Density Functional Theory with London Dispersion Corrections. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2011**, *1*, 211–228.

(27) Mostaghim, S.; Hoffmann, M.; König, P. H.; Frauenheim, T.; Teich, J. Molecular Force Field Parametrization Using Multi-Objective Evolutionary Algorithms. *Proceedings of the 2004 Congress on Evolutionary Computation (IEEE Catal. No.04TH8753)* **2004**, *1*, 212–219.

(28) Jaramillo-Botero, A.; Naserifar, S.; Goddard, W. A. General Multiobjective Force Field Optimization Framework, with Application to Reactive Force Fields for Silicon Carbide. *J. Chem. Theory Comput.* **2014**, *10*, 1426–1439.

(29) Stöbener, K.; Klein, P.; Horsch, M.; Küfer, K.; Hasse, H. Parametrization of Two-Center Lennard-Jones Plus Point-Quadrupole Force Field Models by Multicriteria Optimization. *Fluid Phase Equilif.* **2016**, *411*, 33–42.

(30) Krishnamoorthy, A.; Mishra, A.; Kamal, D.; Hong, S.; Nomura, K.-i.; Tiwari, S.; Nakano, A.; Kalia, R.; Ramprasad, R.; Vashishta, P. EZFF: Python Library for Multi-Objective Parameterization and Uncertainty Quantification of Interatomic Forcefields for Molecular Dynamics. *SoftwareX* **2021**, *13*, 100663.

(31) Krishnamoorthy, A.; Mishra, A.; Grabar, N.; Baradwaj, N.; K. Kalia, R.; Nakano, A.; Vashishta, P. Evolutionary Multi-Objective Optimization and Pareto-Frontal Uncertainty Quantification of Interatomic Forcefields for Thermal Conductivity Simulations. *Comput. Phys. Commun.* **2020**, *254*, 107337.

(32) Miettinen, K. *Nonlinear Multiobjective Optimization*; Springer Science & Business Media: New York, NY, 1998.

(33) Dowling, A. W.; Ruiz-Mercado, G.; Zavala, V. M. A Framework for Multi-Stakeholder Decision-Making and Conflict Resolution. *Comput. Chem. Eng.* **2016**, *90*, 136–150.

(34) 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.

(35) Hülsmann, M.; Reith, D. SpaGrOW—A Derivative-Free Optimization Scheme for Intermolecular Force Field Parameters Based on Sparse Grid Methods. *Entropy* **2013**, *15*, 3640–3687.

(36) Murzyn, K.; Bratek, M.; Pasenkiewicz-Gierula, M. Refined OPLS All-Atom Force Field Parameters for *n*-Pentadecane, Methyl Acetate, and Dimethyl Phosphate. *J. Phys. Chem. B* **2013**, *117*, 16388–16396.

(37) Raabe, G. Molecular Simulation Studies on the Vapor-Liquid Phase Equilibria of Binary Mixtures of R-1234yf and R-1234ze(E) with R-32 and CO<sub>2</sub>. *J. Chem. Eng. Data* **2013**, *58*, 1867–1873.

(38) Zhang, Y.; Zhang, Y.; McCready, M. J.; Maginn, E. J. Evaluation and Refinement of the General AMBER Force Field for Nineteen Pure Organic Electrolyte Solvents. *J. Chem. Eng. Data* **2018**, *63*, 3488–3502.

(39) Lafitte, T.; Apostolakou, A.; Avendano, C.; Galindo, A.; Adjiman, C. S.; Müller, E. A.; Jackson, G. Accurate Statistical Associating Fluid Theory for Chain Molecules Formed From Mie Segments. *J. Chem. Phys.* **2013**, *139*, 154504.

(40) Müller, E. A.; Jackson, G. Force-Field Parameters from the SAFT- $\gamma$  Equation of State for Use in Coarse-Grained Molecular Simulations. *Annu. Rev. Chem. Biomol. Eng.* **2014**, *5*, 405–427.

(41) Rizzi, F.; Najm, H. N.; Debusschere, B. J.; Sargsyan, K.; Salloum, M.; Adalsteinsson, H.; Knio, O. M. Uncertainty Quantification in MD simulations. Part II: Bayesian Inference of Force-Field Parameters. *Multiscale Model. Simul.* **2012**, *10*, 1460–1492.

(42) Zimoń, M. J.; Sawko, R.; Emerson, D. R.; Thompson, C. Uncertainty Quantification at the Molecular–Continuum Model Interface. *Fluids* **2017**, *2*, 12.

(43) Wu, S.; Angelikopoulos, P.; Tauriello, G.; Papadimitriou, C.; Koumoutsakos, P. Fusing Heterogeneous Data for the Calibration of Molecular Dynamics Force Fields Using Hierarchical Bayesian Models. *J. Chem. Phys.* **2016**, *145*, 244112.

(44) Messerly, R. A.; Razavi, S. M.; Shirts, M. R. Configuration-Sampling-Based Surrogate Models for Rapid Parameterization of Non-Bonded Interactions. *J. Chem. Theory Comput.* **2018**, *14*, 3144–3162.

(45) Liu, H.; Fu, Z.; Li, Y.; Sabri, N. F. A.; Bauchy, M. Parameterization of Empirical Forcefields for Glassy Silica Using Machine Learning. *MRS Commun.* **2019**, *9*, 593–599.

(46) Zhang, Y. Computational Methods to Assist in Material Discovery: Membranes and Lithium-Ion Battery Electrolytes. Ph.D. thesis, University of Notre Dame, Notre Dame, IN, 2019.

(47) Pechlaner, M.; Reif, M. M.; Oostenbrink, C. Reparametrization of United-Atom Amine Solvation in the GROMOS Force Field. *Mol. Phys.* **2017**, *115*, 1144–1154.

(48) Diem, M.; Oostenbrink, C. Hamiltonian Reweighting To Refine Protein Backbone Dihedral Angle Parameters in the GROMOS Force Field. *J. Chem. Inf. Model.* **2020**, *60*, 279–288.

(49) Tamura, R.; Lin, J.; Miyazaki, T. Machine Learning Forces Trained by Gaussian Process in Liquid States: Transferability to Temperature and Pressure. *J. Phys. Soc. Jpn.* **2019**, *88*, 044601.

(50) Burn, M. J.; Popelier, P. L. Creating Gaussian Process Regression Models for Molecular Simulations Using Adaptive Sampling. *J. Chem. Phys.* **2020**, *153*, 054111.

(51) Vandermause, J.; Torrisi, S. B.; Batzner, S.; Xie, Y.; Sun, L.; Kolpak, A. M.; Kozinsky, B. On-the-Fly Active Learning of Interpretable Bayesian Force Fields for Atomistic Rare Events. *npj Comput. Mater.* **2020**, *6*, 1–11.

(52) Stephenson, D.; Kermode, J. R.; Lockerby, D. A. Accelerating Multiscale Modelling of Fluids with On-the-Fly Gaussian Process Regression. *Microfluid. Nanofluid.* **2018**, *22*, 1–12.

(53) McDonagh, J. L.; Shkurti, A.; Bray, D. J.; Anderson, R. L.; Pyzer-Knapp, E. O. Utilizing Machine Learning for Efficient Parameterization of Coarse Grained Molecular Force Fields. *J. Chem. Inf. Model.* **2019**, *59*, 4278–4288.

(54) Sestito, J. M.; Thatcher, M. L.; Shu, L.; Harris, T. A.; Wang, Y. Coarse-Grained Force Field Calibration Based on Multiobjective Bayesian Optimization to Simulate Water Diffusion in Poly-*e*-caprolactone. *J. Phys. Chem. A* **2020**, *124*, 5042–5052.

(55) Razi, M.; Narayan, A.; Kirby, R. M.; Bedrov, D. Force-Field Coefficient Optimization of Coarse-Grained Molecular Dynamics Models with a Small Computational Budget. *Comput. Mater. Sci.* **2020**, *176*, 109518.

(56) Boulanger, E.; Huang, L.; Rupakheti, C.; MacKerell, A. D., Jr.; Roux, B. Optimized Lennard-Jones Parameters for Druglike Small Molecules. *J. Chem. Theory Comput.* **2018**, *14*, 3121–3131.

(57) Adorf, C. S.; Dodd, P. M.; Ramasubramani, V.; Glotzer, S. C. Simple Data and Workflow Management with the Signac Framework. *Comput. Mater. Sci.* **2018**, *146*, 220–229.

(58) Ramasubramani, V.; Adorf, C.; Dodd, P.; Dice, B.; Glotzer, S. signac: A Python Framework for Data and Workflow Management. *Proceedings of the Python in Science Conference*; 2018.

(59) Matthews, A. G.; van der Wilk, M.; Nickson, T.; Fujii, K.; Boukouvalas, A.; León-Villagrá, P.; Ghahramani, Z.; Hensman, J. GPflow: A Gaussian Process Library Using TensorFlow. *J. Mach. Learn. Res.* **2017**, *18*, 1–6.

(60) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; Duchesnay, E. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.

(61) Lemmon, E.; Bell, I. H.; Huber, M.; McLinden, M. *NIST Standard Reference Database 23: Reference Fluid Thermodynamic and Transport Properties-REFPROP*, Version 10.0, Standard Reference Data Program, National Institute of Standards and Technology: Gaithersburg, 2018.

(62) Case, D.; Darden, T.; Cheatham, T. E.; Simmerling, L.; Wang, C.; Duke, J.; Luo, R.; Walker, R.; Zhang, R.; Merz, W.; Roberts, K.; Wang, B.; Hayik, B.; Roitberg, S.; Seabra, A.; Kolossváry, G.; Wong, I.; Paesani, K.; Vanicek, F.; Liu, J.; Wu, J.; Brozell, X.; Steinbrecher, S.; Gohlke, T.; Cai, H.; Ye, Q.; Wang, X.; Hsieh, J.; Cui, M.-J.; Roe, G.; Mathews, D.; Sstein, D.; Sagui, M.; Babin, C.; Luchko, V.; Gusarov, T.; Kovalenko, S.; Kollman, A.; AMBER, P. *AMBER 11*; University of California, San Francisco, 2010.

(63) Frisch, M. J.; Trucks, G. W.; Schlegel, H. B.; Scuseria, G. E.; Robb, M. A.; Cheeseman, J. R.; Scalmani, G.; Barone, V.; Petersson, G. A.; Nakatsuji, H.; Li, X.; Caricato, M.; Marenich, A.; Bloino, J.; Janesko, B. G.; Gomperts, R.; Mennucci, B.; Hratchian, H. P.; Ortiz, J. V.; Izmaylov, A. F.; Sonnenberg, J. L.; Williams-Young, D.; Ding, F.; Lipparini, F.; Egidi, F.; Goings, J.; Peng, B.; Petrone, A.; Henderson, T.; Ranasinghe, D.; Zakrzewski, V. G.; Gao, J.; Rega, N.; Zheng, G.; Liang, W.; Hada, M.; Ehara, M.; Toyota, K.; Fukuda, R.; Hasegawa, J.; Ishida, M.; Nakajima, T.; Honda, Y.; Kitao, O.; Nakai, H.; Vreven, T.; Throssell, K.; Montgomery, J. A., Jr.; Peralta, J. E.; Ogliaro, F.; Bearpark, M.; Heyd, J. J.; Brothers, E.; Kudin, K. N.; Staroverov, V. N.; Keith, T.; Kobayashi, R.; Normand, J.; Raghavachari, K.; Rendell, A.; Burant, J. C.; Iyengar, S. S.; Tomasi, J.; Cossi, M.; Millam, J. M.; Klene, M.; Adamo, C.; Cammi, R.; Ochterski, J. W.; Martin, R. L.; Morokuma, K.; Farkas, O.; Foresman, J. B.; Fox, D. J. *Gaussian 09*, Revision E.01; Gaussian, Inc.: Wallingford, CT, 2016.

(64) Becke, A. D. A New Mixing of Hartree-Fock and Local Density-Functional Theories. *J. Chem. Phys.* **1993**, *98*, 1372–1377.

(65) Stephens, P. J.; Devlin, F. J.; Chabalowski, C. F.; Frisch, M. J. *Ab initio* Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields. *J. Phys. Chem.* **1994**, *98*, 11623–11627.

(66) Rasmussen, C. E.; Williams, C. K. *I. Gaussian Processes for Machine Learning*; The MIT Press: Cambridge, MA, 2006.

(67) Bussi, G.; Donadio, D.; Parrinello, M. Canonical Sampling Through Velocity Rescaling. *J. Chem. Phys.* **2007**, *126*, 014101.

(68) Berendsen, H. J. C.; Postma, J. P. M.; van Gunsteren, W. F.; DiNola, A.; Haak, J. R. Molecular Dynamics with Coupling to an External Bath. *J. Chem. Phys.* **1984**, *81*, 3684–3690.

(69) Parrinello, M.; Rahman, A. Polymorphic Transitions in Single Crystals: A New Molecular Dynamics Method. *J. Appl. Phys.* **1981**, *52*, 7182–7190.

(70) Hockney, R. W.; Eastwood, J. W. *Computer Simulation Using Particles*; McGraw-Hill: New York, 1981.

(71) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A Smooth Particle Mesh Ewald Method. *J. Chem. Phys.* **1995**, *103*, 8577–8593.

(72) Hess, B. P-LINCS: A Parallel Linear Constraint Solver for Molecular Simulation. *J. Chem. Theory Comput.* **2008**, *4*, 116–122.

(73) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High Performance Molecular Simulations Through Multi-Level Parallelism From Laptops to Supercomputers. *SoftwareX* **2015**, *1*, 19–25.

(74) DeFever, R. S.; Matsumoto, R. A.; Dowling, A. W.; Cummings, P. T.; Maginn, E. J. MoSDef Cassandra: A Complete Python Interface for the Cassandra Monte Carlo Software. *J. Comput. Chem.* **2021**, *42*, 1321–1331.

(75) Shah, J. K.; Marin-Rimoldi, E.; Mullen, R. G.; Keene, B. P.; Khan, S.; Paluch, A. S.; Rai, N.; Romaniello, L. L.; Rosch, T. W.; Yoo, B.; Maginn, E. J. Cassandra: An Open Source Monte Carlo Package for Molecular Simulation. *J. Comput. Chem.* **2017**, *38*, 1727–1739.

(76) Choi, C. S.; Prask, H. J.; Prince, E. Crystal Structure of  $\text{NH}_4\text{ClO}_4$  at 298, 78, and 10° K by Neutron Diffraction. *J. Chem. Phys.* **1974**, *61*, 3523.

(77) Zhu, W.; Wang, X.; Xiao, J.; Zhu, W.; Sun, H.; Xiao, H. Molecular Dynamics Simulations of AP/HMX Composite with a Modified Force Field. *J. Hazard. Mater.* **2009**, *167*, 810–816.

(78) Tow, G. M.; Maginn, E. J. Evaluating Physical Properties of the Orthorhombic Crystal Phase of Ammonium Perchlorate Using a Class II Force Field. *J. Chem. Phys.* **2018**, *149*, 244502.

(79) Tuckerman, M. E.; Alejandre, J.; López-Rendón, R.; Jochim, A. L.; Martyna, G. J. A Liouville-operator Derived Measure-Preserving Integrator for Molecular Dynamics Simulations in the Isothermal-Isobaric Ensemble. *J. Phys. A: Math. Gen.* **2006**, *39*, 5629–5651.

(80) Shinoda, W.; Shiga, M.; Mikami, M. Rapid Estimation of Elastic Constants by Molecular Dynamics Simulation Under Constant Stress. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2004**, *69*, 134103.

(81) Plimpton, S. Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comput. Phys.* **1995**, *117*, 1–19.

(82) *Handbook for the Montreal Protocol on Substances that Deplete the Ozone Layer*; United Nations Environment Programme, Ozone Secretariat, 2006.

(83) Fermeglia, M.; Ferrone, M.; Pril, S. Development of an All-Atoms Force Field From *ab initio* Calculations for Alternative Refrigerants. *Fluid Phase Equilib.* **2003**, *210*, 105–116.

(84) Lísal, M.; Vacek, V. Effective Potentials for Liquid Simulation of the Alternative Refrigerants HFC-134a ( $\text{CF}_3\text{CH}_2\text{F}$ ) and HFC-125 ( $\text{CF}_3\text{CHF}_2$ ). *Fluid Phase Equilib.* **1997**, *127*, 83–102.

(85) Stoll, J.; Vrabec, J.; Hasse, H. A Set of Molecular Models for Carbon Monoxide and Halogenated Hydrocarbons. *J. Chem. Phys.* **2003**, *119*, 11396–11407.

(86) Cygan, R. T.; Liang, J.-J.; Kalinichev, A. G. Molecular Models of Hydroxide, Oxyhydroxide, and Clay Phases and the Development of a General Force Field. *J. Phys. Chem. B* **2004**, *108*, 1255–1266.

(87) Tow, G. M.; Maginn, E. J. Fully Atomistic Molecular Dynamics Simulations of Hydroxyl-Terminated Polybutadiene with Insights into Hydroxyl Aggregation. *Macromolecules* **2020**, *53*, 2594–2605.

(88) Stöbener, K.; Klein, P.; Reiser, S.; Horsch, M.; Küfer, K.-H.; Hasse, H. Multicriteria Optimization of Molecular Force Fields by Pareto Approach. *Fluid Phase Equilib.* **2014**, *373*, 100–108.

(89) Tenorio, L. *An Introduction to Data Analysis and Uncertainty Quantification for Inverse Problems*; SIAM, 2017.

(90) Madin, O. C.; Boothroyd, S.; Messerly, R. A.; Chodera, J. D.; Fass, J.; Shirts, M. R. Bayesian Inference-Driven Model Parameterization and Model Selection for 2CLJQ Fluid Models. *arXiv.org* **2021**, No. 2105.07863.

(91) Conn, A. R.; Gould, N. I.; Toint, P. L. *Trust Region Methods*; SIAM, 2000.