

Using physical property surrogate models to perform multi-fidelity global optimization of force field parameters

Owen C. Madin¹ and Michael R. Shirts¹

¹Department of Chemical & Biological Engineering, University of Colorado Boulder, Boulder, CO, USA 80309

***For correspondence:**

michael.shirts@colorado.edu (MRS)

Abstract van der Waals dispersion-repulsion interactions, commonly represented in atomistic force fields by the Lennard-Jones (LJ) potential, play an important role in the accuracy of molecular simulations. Training the force field parameters used in the LJ potential is challenging, generally requiring adjustment based on simulations of macroscopic physical properties. The large computational expense of these simulations, especially if many parameters are trained simultaneously, limits the size of training data set and number of optimization steps that can be taken, often requiring modelers to perform optimizations within a local parameter region. To allow for more global LJ parameter optimization against large training sets, we introduce a multi-fidelity optimization technique which uses Gaussian process surrogate modeling to build inexpensive models of physical properties as a function of LJ parameters. This allows for fast evaluation of objective functions, greatly accelerating searches over parameter space and enabling the use of global optimization algorithms. We use an iterative framework which performs optimization with differential evolution at the surrogate level, followed by validation at the simulation level and surrogate refinement. Using this technique on two previously studied training sets, containing up to 195 physical property targets, we refit a subset of the LJ parameters for the OpenFF 1.0.0 (Parsley) force field. We demonstrate that this multi-fidelity technique can find improved parameter sets compared to a purely simulation-based optimization by searching more broadly and escaping local minima. In particular, this technique often finds significantly different parameter minima that have comparably accurate performance. In most cases, these parameter sets are transferable to other similar molecules in a test set. This multi-fidelity technique provides a platform for fast optimization against physical properties that can be refined and applied in multiple ways to the development of molecular models.

1 Introduction

1.1 Accurate force fields are important in computational biophysics

Accurate molecular interaction potentials, usually referred to as force fields, are an essential part of modern molecular dynamics workflows. For common applications such as simulations of proteins and computer aided drug design (CADD), the simple fixed-charge force field functional form [1–3] is generally used. This formulation splits the potential energy of molecules into discrete components, with separate energy terms for each component [4]. Broadly, these can be divided into the bonded (or valence) components, which comprise the energies of bond lengths, bond angles, and torsional angles, and the non-bonded components, representing short-range dispersion repulsion interactions and longer-range Coulombic interactions.

This type of force field has been successful in many applications because of its simplicity, interpretability, and computational efficiency. Many studies have used these force fields to probe the mechanisms of protein dynamics [5–8], and they have become widely adopted in the pharmaceutical industry as a means of screening drug candidate molecules *in silico* [8–12]. While these force fields are quite simple in their functional form, their accuracy is dependent on hundreds to thousands of parameters, which dictate the strength of interactions in different chemical environments. Decades of effort from the computational chemistry community has produced many different parameter sets to cover a wide range of chemistries [3, 13–16], largely by fitting parameters to quantum mechanics (QM) calculations [17, 18] and experimental physical properties [19, 20].

1.2 Non-bonded training is expensive and difficult

Over the years, fitting of the bonded parameters has perhaps received the most attention, due to their importance in determining the internal structure of molecules and proteins, and the relative ease of generating QM data for the molecules of interest. Fitting the atomic partial charges used in the Coulombic potential has also received significant attention, but is slightly more difficult, as mapping an continuous electrostatic potential onto a set of discrete atoms is conformation-dependent and involves a loss of fidelity. However, modelers have achieved good results using QM-based methods such as RESP [21, 22] and semi-empirical methods such as AM1-BCC [23, 24].

The dispersion-repulsion interactions, almost always modeled with the Lennard-Jones (LJ) potential, have received the least attention in fitting, as they are typically trained against experimental physical property data [19, 20, 25], since obtaining dispersion-repulsion estimates from QM is difficult [26?]. This leads to challenges with curating appropriate sets of experimental physical property data from the literature, as well as the computational cost of simulating sets of physical property data with molecular dynamics. Most physical properties used in training, which include densities [19], enthalpies of vaporization [19], enthalpies of mixing [27], solvation free energies [20] and dielectric constants [28, 29], require equilibrium simulations in one or more phases, and in some cases may require alchemical simulation techniques [30]. In conjunction with the need to train against larger datasets to ensure accuracy and transferability, this makes optimization of LJ parameters a difficult problem. Calculating a single objective function value in order to measure parameter fitness requires a large number of simulations, which can be difficult to coordinate and execute, especially depending on available computational resources. As a result, one can find many instances of LJ parameters in major force fields that have remained unchanged for more than 20 years, despite significant advancements in hardware, simulation software, and methodology in that time.

Recently, as part of the Open Force Field (OpenFF) Initiative, we have examined new methods of LJ parameter optimization. Central to these efforts is the development of the OpenFF Evaluator simulation workflow driver [31], which provides a standardized set of workflows for automatically building and executing physical property simulations for a given training or test data set. With the automation that this software provides, we can apply parameter optimization techniques to LJ parameters with minimal human intervention. In particular, this software enabled the application of the ForceBalance [32] parameter optimization package to LJ parameters. Using regularized least squares optimization with the L-BFGS-B algorithm [33], we minimized an objective function that captures the ability of a parameter set to reproduce physical property observables. Using this framework, we studied the benefits of including physical property data of mixtures in training LJ parameters [27], then applied that training method to a production force field, OpenFF 2.0.0 (known also as Sage)[34].

While we have achieved improved parameter sets with simulation-based regularized least-squares optimization, it has significant limitations. A major drawback is that, regardless of the optimization algorithm used, the number of objective function evaluations possible is limited by the computational cost of simulations. This limits the number of parameter sets that can be considered during optimization, making it difficult to explore high-dimensional and complex parameter spaces. This also necessitates the use of cheaper local optimization methods such as L-BFGS-B with termination after a set number of steps [27]. When coupled with a regularization term in the objective function, included both to ensure the stability of the optimization and to guard against overfitting [32], local optimization methods have a high probability

of remaining in any local minima dictated by its initial values. This means that our ability to explore new areas of parameter space that may provide significant improvement is blunted because of the expense of evaluating the objective.

1.3 Surrogate modeling can accelerate non-bonded training

To facilitate faster evaluation of complex objective functions, modelers often use surrogate models [35, 36], which are meant to approximate an expensive function with a simpler alternative that captures most of the important information. Surrogate modeling techniques have been developed in response to diverse sets of scientific and engineering challenges, including geological modeling [37, 38], engineering design [39, 40], and chemical process modeling [41, 42]. A popular technique for surrogate modeling is Gaussian process (GP) surrogate modeling, which has seen adoption in many disciplines [43–45] due to its simplicity and efficacy in data-sparse regimes. While these surrogates are not perfect imitations of the high-level response, we can construct them to a reasonable level of accuracy with a limited number of expensive evaluations. In the context of molecular simulation parameter optimization, Befort et al. [46] demonstrated a method of optimizing LJ parameters by building GP surrogates based on physical properties and applied this method to several hydrofluorocarbons as well as ammonium perchlorate.

In this paper, we build on this approach, as well as engineering optimization literature [47] and our OpenFF Evaluator software, to introduce a multi-fidelity optimization framework based on the construction of Gaussian process (GP) surrogate models that simultaneously approximate the response surface of many physical properties with respect to changes in the LJ parameters. Using the accelerated objective evaluation offered by the surrogates, we implement a global optimization algorithm to search broadly and propose candidate parameter sets. We then validate these parameter sets by evaluating the objective at the simulation level, accepting candidates in good agreement. Iterating between global optimization over the surrogate, and simulation-level validation and surrogate refinement, we can drive the optimizer to explore promising regions of parameter space with a limited number of simulation evaluations.

We test this approach by performing multi-fidelity LJ optimization for 12 LJ parameters from the OpenFF 1.0.0 (Parsley) force field. In training, we use a set of 56 pure compound physical properties curated in a previous paper [27]. We also benchmark the results against test sets curated in the same paper; while newer versions of OpenFF exist, using OpenFF 1.0.0 allows a direct comparison to the results of our previous optimization. With this optimization problem, we characterize the optimization method, discussing reproducibility, seed configurations, and optimization trajectories. We also show that this method can be extended to larger problems, by applying it to a larger training set of 195 mixture properties from the same paper.

2 Methods

2.1 Optimization Strategy

Our optimization strategy aims to minimize an objective function $\chi(\theta)$ as in Equation 1, where θ is a vector of force field parameters, and χ is some measure of the quality of those parameters.

$$\min_{\theta} \chi(\theta) \quad (1)$$

In our applications, parameter fitness is described by the ability of a force field containing those parameters to reproduce a specific training set of experimental physical properties, although we note that this strategy could be also be applied to a training set containing quantities from QM simulations. The optimization strategy we employ is adapted from the framework proposed by Dennis and Troczon [47] and features two levels of fidelity for estimating the objective function for a parameter set:

- “Simulation level”, where the objective function is directly evaluated by using molecular dynamics to simulate the training set with a force field containing the parameter set. This is considered to be the “ground truth”, as it most accurately represents the force field’s performance, although there is some level of statistical uncertainty due to the stochasticity in the molecular dynamics simulation.

- “Surrogate level”, where the objective function is estimated by a collection of surrogate models that approximate the result of a simulation level evaluation of the training set. The surrogate level evaluation of the objective function has systematic uncertainty where the surrogates approximation deviates from the simulation-level estimation of the training set. It may also have statistical uncertainty depending on the type of surrogate used, although the surrogates that we use do not.

Our optimization strategy relies on the cheaper but less accurate surrogate level to perform most of the optimization, using the more accurate and expensive simulation level only to build the surrogates and validates proposed surrogate-level solutions. The optimization framework is illustrated in Figure 1.

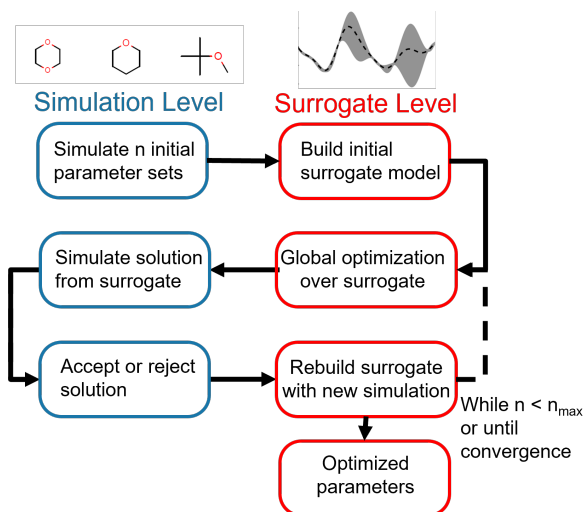


Figure 1. Flowchart of multi-fidelity optimization strategy. Optimization is initialized by simulating an initial sample of parameter vectors. Surrogate models for each of the physical properties in the training set are then built from this initial sample. Global optimization is then performed at the surrogate level, utilizing the speedup gained with surrogate-level fast objective evaluation. Once this proposes a candidate vector of optimized parameters, the objective function for that parameter vector is evaluated at the simulation level. If the simulation-level objective is lower than the simulation objective for the previous parameter vector, the new parameter vector is accepted as an improved solution; if not, it is rejected. Regardless of acceptance or rejection, the surrogate model is rebuilt with the information from the simulation level evaluation. This process is then repeated until a maximum number of simulation optimizations is reached, a convergence criteria is met, or the optimizer cannot find an improved solution.

The advantage of this strategy is its use of the properties of both the surrogate and simulation level to drive optimization. While surrogate level evaluation is much faster than simulation level evaluation, surrogates need simulation points in the region of interest to accurately reproduce the objective function. Since parameter spaces are large and the region of interest is not known *a priori*, an exhaustive strategy would require a very large number of simulation level evaluations to build globally accurate surrogates, negating the speedup gained by using surrogates. We instead build a surrogate from a minimal initial set of evaluations of the objective function and allow the surrogate to suggest new parameter vectors for the simulation level to evaluate. We can therefore iteratively drive the optimization towards the region of interest without incurring too much computational cost, acquiring more information to improve the surrogates along the way. This allows us to pair the global optimization strategies available at the surrogate level with the accuracy of simulation-level validation.

This strategy is sufficiently general to allow for the use of a large variety of objective functions, global optimization techniques, and surrogate modeling strategies. In this particular study we focus on only single combination: a weighted least-squares objective function, the differential evolution global optimization algorithm, and Gaussian Process (GP) surrogate modeling.

2.1.1 Objective Function

The objective function we use, shown in Equation 2, is adapted from the type used in the ForceBalance optimization software package and used in our previous work [27]. While that objective included a regularization term for stability and to prevent overfitting, we omit that term, allowing our algorithm to search more broadly to find optimized parameter vectors.

$$\chi(\theta) = \sum_{n=1}^N \frac{1}{M_n} \sum_{m=1}^{M_n} \left(\frac{y_m^{ref} - y_m(\theta)}{d_n} \right)^2 \quad (2)$$

In this equation, we consider N types of physical properties, each with some number M_n of measurements for that type. The quantity y_m represents the value of the m th measurement for a physical property type, and the denominator d_n is a scaling coefficient for a given property. The scaling coefficients are set so that each physical property type contributes equally to the objective function for OpenFF 1.0.0 [18], the starting point of our optimizations.

2.1.2 Global Optimization

We use the differential evolution [48] global optimization algorithm, as implemented in the scipy Python package version 1.7.0 [49]. Differential evolution is a stochastic direct search algorithm similar to other genetic algorithm strategies [50]. In this strategy, a set of N initial vectors is proposed, and then “mutated” by randomly increasing or decreasing elements of the vector, and recombined, by randomly replacing some elements of the vector with elements of other vectors. The objective function is then evaluated for each of the proposed vectors, and a new set of vectors is proposed based on the lowest objective functions. This process is repeated until convergence, when new vectors no longer outperform the current solutions. We use the default optimization parameters in scipy, with a population size of 180 vectors, an iteration-dependent mutation constant selected from the range (0.5,1), and a recombination constant of 0.7. We note that, depending on the problem, a single iteration of this algorithm requires between $10^3 - 10^4$ objective function evaluations *per iteration*, each of which depends on the value of 50-200 physical properties.

The bounds of the global optimization over the surrogate model are determined by the parameter sets used to build the surrogate. For each parameter θ_i in the parameter vector θ , the bounds for θ_i are determined by the minimum ($\min(\theta_i)$) and maximum value ($\max(\theta_i)$) of θ_i in the set of parameter vectors $\Theta = [\theta^1, \theta^2, \dots, \theta^N]$ used to build the surrogate. We then apply a small expansion factor η to the parameter range, so that the optimization algorithm can search outside of the initial simulation box (described in Section 2.4.2). To expand the box, we multiply $\max(\theta_i)$ by η , and divide $\min(\theta_i)$ by η to form the bounds box $[\text{LB}(\theta_i), \text{UB}(\theta_i)]$. We chose the value of η to be 1.1, expanding the box 10% in each direction, to allow the optimizer to search aggressively. This process is repeated for each parameter $\theta_i \in \theta$ to form an N -dimensional box, and is described in Equation 3.

$$\text{LB}(\theta_i) = \frac{1}{\eta} \times \min_{\theta \in \Theta}(\theta_i), \text{UB}(\theta_i) = \eta \times \max_{\theta \in \Theta}(\theta_i) \quad (3)$$

2.2 Construction of Physical Property Surrogates

GP surrogate models are built with the botorch [51] software package, version 0.6.0, which provides a convenient and extensible framework for building a large number of surrogates. Surrogate are constructed individually for each physical property in the test set, and objective functions are calculated from the surrogates’ predictions of their respective physical properties; the surrogate does not predict the objective function directly. Surrogates are constructed from all of the simulation level evaluations available; e.g. if there are simulations of 20 physical properties with 10 different LJ parameters vectors, then our process builds 20 individual surrogate models, each using all 10 parameter vectors in their construction. All surrogates use a constant mean function and RBF (radial basis function) covariance kernel; independent length scales l for each parameter are chosen using automatic relevance determination (ARD) [52]. The length scales l represent the variances in each parameter dimension of the radial basis function; essentially, the diagonal of the covariance matrix; all off-diagonal elements (covariances) are set to 0. If a physical property

simulated at a given set of parameters finishes with errors, that set is omitted from surrogate building; additionally, any sets that have density measurements lower than 20% of the experimental value are omitted, as this likely indicates that the parameters have induced a phase change. However, this restriction did not affect the optimization, as parameters violating the density constraint were never produced through our optimizations.

In cases where an optimization over a surrogate model failed to find a lower objective value than the current simulation objective, surrogates are rebuilt with constraints on the length scales l used for the variances of each parameter. This approach was chosen as we found in testing that optimizations may fail because a surrogate was set with a length scale too low during ARD, producing a surrogate with poor quality. If the optimizer cannot find a better objective value over the surrogate, it is first rebuilt with a length scale constraint such that $l > 10^{-10}$; if this is not successful, a stricter length scale constraint of $l > 10^{-5}$ is imposed. If this is still not successful, the optimization is terminated. The quality of the surrogate model is reduced when length scale constraints are introduced, so constraints are not used unless an optimization fails.

2.3 Physical Property Simulations

Physical property simulations were handled with the OpenFF Evaluator[31] software package, version 0.3.4 [53], using the default workflows [54] for all properties simulated. We performed simulations to estimate pure density (ρ_L), mixture density ($\rho_L(x)$), enthalpy of vaporization (ΔH_{vap}) and enthalpy of mixing ($\Delta H_{mix}(x)$). We perform all condensed-phase simulations in the NPT ensemble, with initial simulation boxes of 1000 molecules built using PackMOL [55]. After building the boxes, we perform an energy minimization on the simulation boxes, followed by a 0.2 ns equilibration simulation and a 2 ns production simulation, which is sufficient to converge these simple physical properties[31]. In the calculation of ΔH_{vap} , we use 30 ns single molecular NVT simulations without periodic boundary conditions to estimate the gas phase energies. All simulations use a 2 fs timestep and a Langevin integrator with BAOAB splitting [56]. More complete simulation details are available in our previous work [27], which uses the same simulation workflows.

2.4 Optimization tasks

We focused on two separate optimization tasks, both initially developed in our previous study [27]. Both tasks optimize the same set of 12 LJ parameters ($R_{min/2}$ and ϵ for 6 LJ SMIRKS types), and both use the same small molecules (alkanes, alcohols, ethers, esters and ketones) in their training sets. The tasks are differentiated by the different types of physical property training data that are used in the evaluation of the objective function:

1. "Pure only": This task focuses on optimizing the LJ parameters against a set of 56 pure compound measurements, ρ_L and ΔH_{vap} for each of 28 compounds in the training set. We used this task to test the optimization strategy, as it represents the typical type of training set used in LJ optimization, and has relatively low computational expense because of the number and types of physical properties that need to be estimated.
2. "Mixture only": This task optimizes the LJ parameters against a larger set of 195 physical properties of binary mixtures, which include 195 total measurements of ΔH_{mix} and $\rho_L(x)$. This task extends the strategy to a significantly larger training set, and represents the type of training set that we determined to perform best in our previous study.

While we reported optimized parameter sets for these tasks in our previous work, here we use those results as a baseline to test our multi-fidelity strategy.

2.4.1 Parameters to be optimized

We optimize the LJ $R_{min/2}$ and ϵ for 6 LJ types, which are described in Table 1. We also note that several LJ types are exercised by molecules in the training set, but are not optimized, due to either having very specific chemical contexts that are not exercised widely enough to optimize, or, in the case of the [#1:1] - [#8] (hydroxyl hydrogen) parameter, because the ϵ has been set to an arbitrary small non-zero value to avoid unphysical effects [57].

Refit SMIRKS Type	Description
Refit parameters	
[#1:1] - [#6X4]	Hydrogen attached to tetravalent carbon
[#6:1]	Generic carbon
[#6X4:1]	Tetravalent carbon
[#8:1]	Generic oxygen
[#8X2H0+0:1]	Divalent oxygen with no hydrogens attached
[#8X2H1+0:1]	Divalent oxygen with one hydrogen attached
Parameters exercised but not optimized	
[#1:1] - ([#6X4]) - [#7, #8, #9, #16, #17, #35]	Hydrogen attached to tetravalent carbon attached to an electronegative atom
[#1:1] - [#6X3] (~[#7, #8, #9, #16, #17, #35]) ~[#7, #8, #9, #16, #17, #35]	Hydrogen attached to trivalent carbon attached to 2 electronegative atoms
[#1:1] - [#8]	Hydrogen attached to oxygen

Table 1. All LJ SMIRKS types adjusted in the training of OpenFF 2.0.0, along with descriptions of the chemical contexts they describe. LJ ϵ and $R_{min/2}$ are adjusted for each of these types.

2.4.2 Initial physical property simulations

To build an initial surrogate in each optimization, we simulate an initial set of N parameter vectors, one of which is always the parameter vector corresponding to OpenFF 1.0.0. We select these vectors from an initial parameter space, described in Table 2. The size of this space is measured in percentage of the parameter values from OpenFF 1.0.0, and is determined from the results of our previous optimization study, based on how much each parameter was adjusted in that study. From this space, we select the remaining $N - 1$ parameter vectors using Latin Hypercube Sampling (LHS), as implemented with the Surrogate Modeling Toolbox [58] (smt) Python, library version 1.1.0.

Refit SMIRKS Type	ϵ initial parameter range (% of OpenFF 1.0.0)	$R_{min/2}$ initial parameter range (% of OpenFF 1.0.0)
Refit parameters		
[#1:1] - [#6X4]	(50%, 150%)	(95%, 105%)
[#6:1]	(90%, 110%)	(95%, 105%)
[#6X4:1]	(90%, 110%)	(95%, 105%)
[#8:1]	(95%, 105%)	(95%, 105%)
[#8X2H0+0:1]	(95%, 105%)	(95%, 105%)
[#8X2H1+0:1]	(95%, 105%)	(95%, 105%)

Table 2. All LJ SMIRKS types adjusted in the training of OpenFF 2.0.0, along with descriptions of the chemical contexts they describe. LJ ϵ and $R_{min/2}$ are adjusted for each of these types.

2.4.3 "Pure only" optimization task

The "pure only" optimization task optimizes the LJ parameters against 56 physical properties (ρ_L , ΔH_{vap}) for a set of 28 molecules, which are shown in Figure 2. A list of the molecules in the "pure only" training set are available in the Supporting Information, Section S1.1.




Figure 2. Molecules in the “pure only” training set. Physical properties in this set include one measurement each of ρ_L and ΔH_{vap} , and are sourced from either the NIST ThermoML Archive (ρ_L) or hand-curated from literature (ΔH_{vap}).

The measurements here are either sourced from the NIST ThermoML Archive [59, 60] (ρ_L) or hand-curated from literature (ΔH_{vap}) [61–72], because of the low number of ΔH_{vap} data points in the ThermoML Archive. All measurements are selected at temperatures and pressures close to ambient (~ 1 atm, 273.15–318.15 K).

For this optimization task, we performed optimizations using $N = 5$ and $N = 10$ initial points, in order to test the effect of the number of initial simulation points on the performance of the algorithm. We performed 5 replicates for both $N = 5$ and $N = 10$ initial points, in order to assess the consistency of the algorithm. For the $N = 10$ replicates, a different set of 9 LHS initial points is selected each time; for the $N = 5$ replicates, each set of initial points is formed by subsampling 4 LHS points from one of the $N = 10$ replicate initial sets, in order to minimize simulation expense.

2.4.4 “Mixture only” training set

The “mixture only” optimization task optimizes the LJ parameters against a set of 195 physical properties ($\rho_L(x)$, $\Delta H_{mix}(x)$) for the set of molecular pairs shown in Figure 3. These molecule pairs are drawn from the same set of molecules as used in the “pure only” training set. All measurements here are selected from the NIST ThermoML archive, and are selected at temperatures and pressures close to ambient (~ 1 atm, 273.15–318.15 K). We select measurements at concentrations with 0.05 mole fraction of 3 target concentrations for each mixture, where available: ($x_1 = 0.25, x_2 = 0.75$), ($x_1 = 0.5, x_2 = 0.5$), ($x_1 = 0.75, x_2 = 0.25$). If no measurements are available within 0.05 mole fraction of any target concentration, no data point is selected for that target concentration. A list of the mixtures in the “mixture only” training set is available in the Supporting Information, Section S1.2.




Figure 3. Molecules in the “mixture only” training set. Physical properties in this set include measurements of $\rho_L(x)$ and $\Delta H_{mix}(x)$ at conditions close to ambient (~ 1 atm, 273.15–318.15 K), and several concentrations ($x_1 = 0.25, x_2 = 0.75$, $x_1 = 0.5, x_2 = 0.5$, $x_1 = 0.75, x_2 = 0.25$), where available, yielding a total of 195 measurements. All measurements are sourced from the NIST ThermoML Archive.

For this optimization task, we performed an initial optimization using $N = 20$ initial points, due to the increased complexity of the training set. After this optimization, we performed a second optimization using $N = 10$ initial points, in order to test whether a more data-sparse optimization could be successful. For the $N = 10$ replicate, 9 initial points are subsampled from the 19 LHS points used in the $N = 20$ replicate to minimize simulation expense.

2.5 Benchmarking

To assess the quality and transferability of the parameter sets produced by our optimization, we tested them on a set of physical properties (29 ρ_L , 318 $\rho_L(x)$, 29 ΔH_{vap} , and 236 $\Delta H_{mix}(x)$) for a new set of molecules and molecular pairs, which serves as the test set. This data set was curated for our previous work, and its selection and composition are discussed there. Physical properties in this set are either hand-curated from literature (ρ_L , ΔH_{vap}), or are selected automatically from the NIST ThermoML Archive. Benchmarking simulations are performed using the same OpenFF Evaluator workflows as simulations used in the optimization process.

3 Results & Discussion

3.1 Pure training set

3.1.1 Optimization

Optimization was generally successful with $N = 5$ and $N = 10$ initial parameter vectors, as it reached significantly lower objective function values than the initial force field in every case. Additionally, when comparing training set RMSE for ΔH_{vap} , all optimization replicates significantly outperform the regularized least squares optimization.

Out of the 10 optimizations run, 4 of them terminated early after the surrogate optimizer could not find an improved solution. This is related to the issues with ARD noted in section 2.2. This suggests that further research is needed to improve the robustness of the surrogate model. Optimizations used between 15–25 total simulations, compared to the 12 used in the simulation-only optimization.

The optimization trajectories of the replicates starting from $N = 5$ initial points are shown in Figure 4, with the evaluated objective function of the previous simulation-only regularized least-squares optimization for comparison.




Figure 4. Performance of the multi-fidelity optimization algorithm on the “pure only” training set, for replicates run with $N = 5$ initial parameter vectors. Left panel shows the objective function at each iteration of the optimization. Right two panels show the RMSE for ΔH_{mix} and $\rho_L(x)$ for the two optimizations, as well as OpenFF 1.0.0 and the previously simulation-only optimization (labeled “sim only” in the graphs). Error bars represent 95% confidence intervals, computed with bootstrapping over the set of molecules in the training set.

We see that in most cases, the optimizer struggles initially, with a high percentage of solutions proposed and rejected in the first 8 steps. While these steps do not immediately yield an improved force field, the parameter vectors they propose are added to the pool of parameter vectors used to build surrogates, eventually “exploring” enough space to find an improved solution, with an average objective of 0.039 vs. an initial

objective of 0.16, an average ρ_L training set RMSE of 0.016 g/mL (initial RMSE 0.027 g/mL), and an average ΔH_{vap} RMSE of 2.75 kJ/mol (initial RMSE 7.15 kJ/mol).

In order to find these improved solutions, the optimization algorithm searches widely and finds a number of qualitatively distinct minima. The optimization trajectories, as well as the optimization trajectory from the simulation-only optimization against the same training set, are shown in Figure 5.

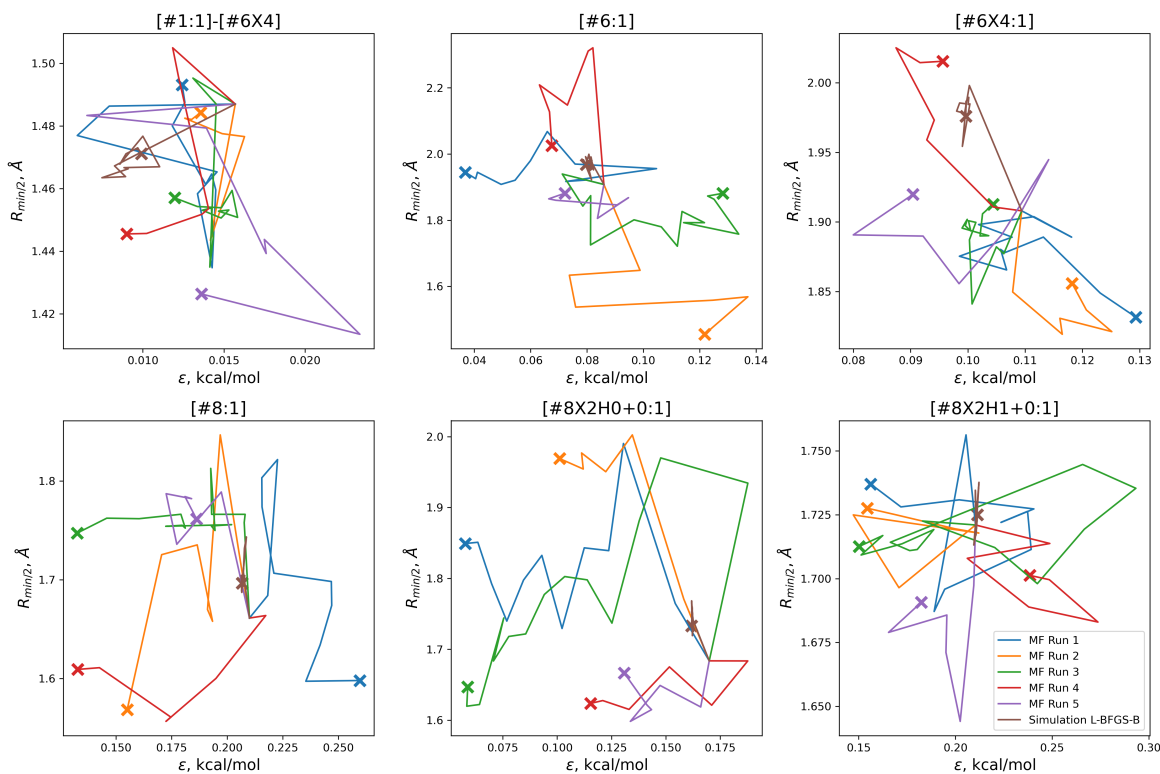


Figure 5. Optimization trajectories for each of the replicates run with 5 initial parameter vectors, as well as the trajectory for the previous simulation-only optimization (brown). Trajectories show that our optimization technique searches widely and finds many separate solutions for this optimization problem.

In comparison to the simulation-only optimization, shown in brown, the replicates of our multi-fidelity optimization searches the parameter space much more broadly. Particularly, the values of oxygen and carbon parameters stay within a narrow range in the simulation-only optimization, but vary widely with our multi-fidelity technique.

The training set RMSEs and objective functions for the $N = 10$ runs are shown in Figure 6.




Figure 6. Performance of the optimization algorithm on the “pure only” training set, for replicates run with $N = 10$ initial parameter vectors. Left panel shows the objective function at each iteration of the optimization. Right two panels show the RMSE for ΔH_{mix} and $\rho_L(x)$ for the two optimizations, as well as OpenFF 1.0.0 and the previous simulation-only optimization (labeled “sim only” in the graphs). Error bars represent 95% confidence intervals, computed with bootstrapping over the set of molecules in the training set.

We note that the optimization initialized with $N = 10$ initial parameter vectors improve the objective function with fewer iterations than the $N = 5$ optimizations. These optimizations leverage the additional initial information to find improved parameter sets sooner, with the trade-off that more simulations are required upfront to achieve this result. The effectiveness of the optimization is also slightly improved over the $N = 5$ optimizations, with an average objective function of 0.031 ($N = 5$: 0.039), an average ρ_L training set RMSE of 0.014 g/mL ($N = 5$: 0.016), and an average ΔH_{vap} RMSE of 2.38 kJ/mol ($N = 5$: 2.75 kJ/mol).

Interestingly, the parameter trajectories in the $N = 10$ optimizations explore a wider range of parameters than the $N = 5$ optimizations, as shown in Figure 7. One might expect that an optimizer with less information would take larger jumps as it attempts to find a minima, but it is likely here that optimization with a less specified initial surrogate ($N = 5$) proposes more solutions that are rejected and therefore finds more conservative minima.

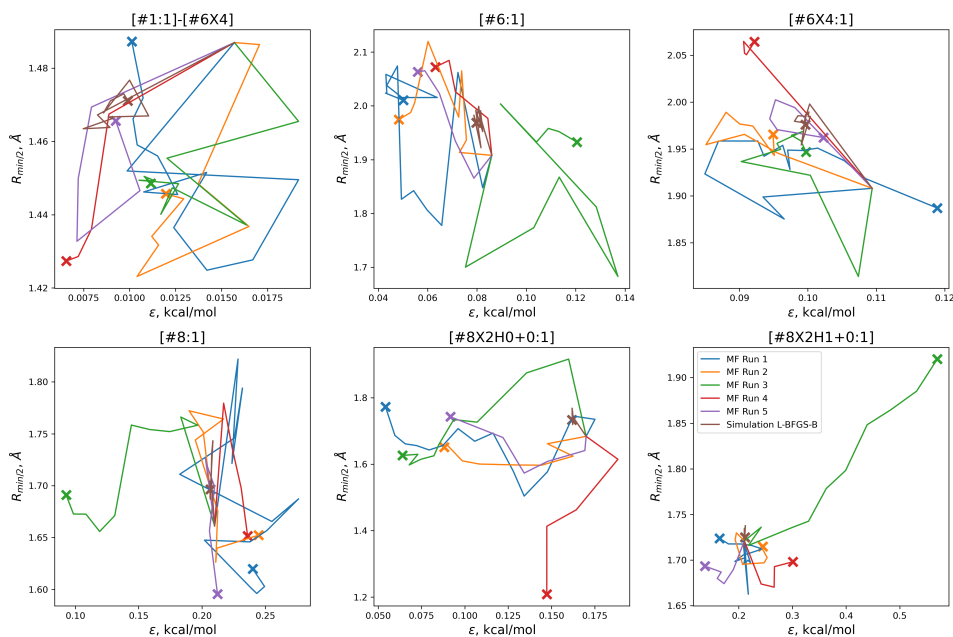


Figure 7. Optimization trajectories for each of the replicates run with 10 initial parameter vectors. Trajectories show that our $N = 10$ initial parameter vector optimization technique searches more widely than our $N = 5$ technique, but is also more likely to drastically alter a parameter while deeply “exploring” a promising candidate, as shown for runs 3 and 4.

Particularly in the oxygen parameters for runs 3 and 4, we can see that this optimization setting can make some very large parameter changes; run 3 has the lowest overall objective, but more than doubles the hydroxyl oxygen ϵ . This may suggest that adding some regularization could benefit the transferability of the optimization, but it also reflects that the ratio of optimization targets to inputs (56:12) leads to an optimization where many solutions can be found.

3.1.2 Parameter Interpretation

Since the optimizations find a large number of separate and diverse solutions, and the set of parameters is small enough to be reasonably interpretable, it is worth examining some of the parameter changes to understand their physical basis and inform future optimizations. Here we analyze some of the most notable changes from the $N = 10$ replicates, which had larger parameter changes relative to the $N = 5$ replicates. The change in parameters from the original OpenFF 1.0.0 is shown in Figure 8.




Figure 8. Changes in parameter values after optimization against the “pure only” data set, relative to OpenFF 1.0.0 (the initial values of the optimization), for the $N = 10$ optimization replicates, as well as the simulation-only solution previously obtained. Parameter changes are typically larger in multi-fidelity optimization compared to the simulation-only optimization, indicating improved exploration of the parameter space; however, this also leads to outliers (hydroxyl oxygen ϵ , ether oxygen $R_{min}/2$).

To identify what points in the training dataset the parameter changes are affecting, we examine the bias of the physical properties training dataset, as measured by mean signed error (MSE). The bias before and after training, as a function of chemical group, is plotted for multi-fidelity run 1, shown in Figure 9. Run 1 is shown as an example as it has one of the best objective functions and no unusual parameter changes (such as occurred in runs 3 and 4), and MSE values are similar for all optimization runs. Similar plots of RMSE and MSE for all 5 runs are available in the Supporting Information, Section S2.2.

N=10, replicate 1

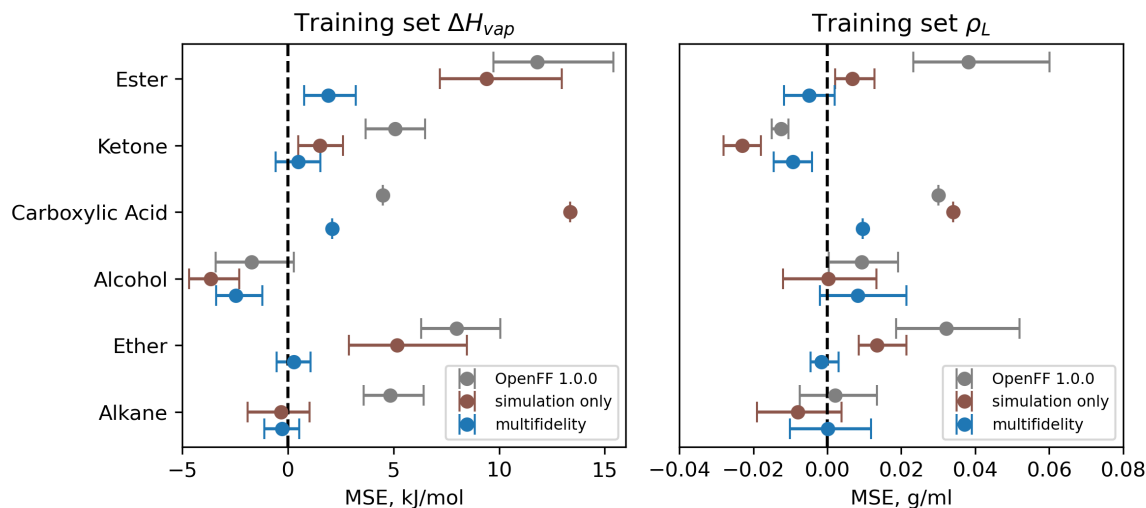


Figure 9. Bias in training set ΔH_{vap} and ρ_L by chemical functionality, as measured by the mean signed error (MSE), for OpenFF 1.0.0 and retrained parameters from $N = 10$ multi-fidelity run 1. Training with multi-fidelity optimization reduces or eliminates bias in several chemical functionalities, including alkanes, ketones, and ethers, where MSE is near 0 kJ/mol. Reduction in MSE shown for run 1 is typical of all “pure only” $N = 10$ multi-fidelity runs; training set MSEs and RMSEs for all multi-fidelity optimization are available in Supporting Information, Section 2.1-2.3. Error bars represent bootstrapped 95% confidence intervals

One consistent trend in most optimizations is the significant overall reduction of ϵ . By decreasing the ϵ s, cohesive forces in the liquid phase are reduced, lowering the barrier for “liberating” a molecule from the gas phase and thereby lowering the enthalpy of vaporization. In OpenFF 1.0.0, the enthalpy of vaporization measurements in the training set have a mean signed error (MSE) of 5.03 kJ/mol, with all moieties except alcohols having a positive deviation from experiment. After multi-fidelity optimization, the training sets have an average MSE (across all multi-fidelity runs) of 0.22 kJ/mol.

The trend of reduced ϵ 's is strongest for the [#1:1]-[#6X4] (hydrogen attached to tetravalent carbon) and [#8X2H0+0:1] (divalent oxygen with 0 hydrogens attached) atom types. The [#1:1]-[#6X4] type is exercised in all molecules in the training set, so reducing the ϵ for this type helps to reduce this overall bias. For the alkanes in the set, [#1:1]-[#6X4] is one of two parameters exercised (along with [#6X4:1]), and alkane ΔH_{vap} training set MSE is reduced from 4.75 kJ/mol in OpenFF 1.0.0 to an average of -0.01 kJ/mol, virtually eliminating the error. The [#8X2H0+0:1] type is exclusive to ether oxygens in this training set, and reducing its ϵ helps to correct a significant overprediction of ether ΔH_{vap} in OpenFF 1.0.0, reducing the ether MSE from 7.95 kJ/mol to an average value of 0.54 kJ/mol after training. This reduction is much larger than the reduction observed after simulation-only optimization.

The ϵ 's for [#6:1] and [#8:1] present an interesting case in multidimensional optimization. We see significant changes in the ϵ 's for both [#6:1] and [#8:1], which represent generic hydrogens and oxygens. However, the presence of more specific types in the training set means that these two types are only exercised together in a C=O double bond (a ketone, ester, or carboxylic acid). In all optimizations but run 3, we see a large reduction in [#6:1] ϵ and a slight increase in [#8:1] ϵ . The adjustment of these parameters, along with an increase in the [#6:1] $R_{min/2}$, corrects an overprediction in the ester and ketone ΔH_{vap} . Notably, simulation-only optimization against the same training set was not able to achieve the same correction for esters.

Interestingly, run 3 takes an opposite approach, increasing ϵ for [#6:1] and decreasing ϵ for [#8:1] but achieving a similar reduction in bias. This suggests that, for the purpose of this optimization, [#6:1] and [#8:1] are treated as a unit. This is not desirable in a larger context, as these parameters are not inherently coupled and will probably lead to significant errors if used in different contexts.

For $R_{min/2}$, the most consistent changes are in [#6:1] and [#6X4:1], which are generally increased. Overall, the effect of increasing $R_{min/2}$ should be to decrease density, as it increases inter-atomic distances and leads to higher molecular volume. Overall, densities are slightly overpredicted in OpenFF 1.0.0, but those overpredictions are concentrated in ethers and esters. For [#6:1], the increase in $R_{min/2}$ helps to reduce a significant overprediction of ester densities.

3.1.3 Surrogate Analysis

Given that the optimizations used in this process vary widely in the sets of parameters that they produce, it is useful to characterize the quality of the surrogate models over the parameter space. Specifically, we measure their global accuracy by comparing the generated surrogate models. We also measure the roughness over the surrogate models by running multiple minimizations on the final surrogate models. We performed this analysis for the $N = 10$ optimization runs.

To assess the ability of the surrogate to make accurate predictions outside its the region of the minimization, we calculated objective functions with the surrogate produced in each $N = 10$ optimization for the parameter sets produced from all other $N = 10$ optimizations. If the surrogates were globally predictive, we would observe low prediction error for the other minima; with high error, surrogates are likely only locally predictive. Results are shown in Figure 10.

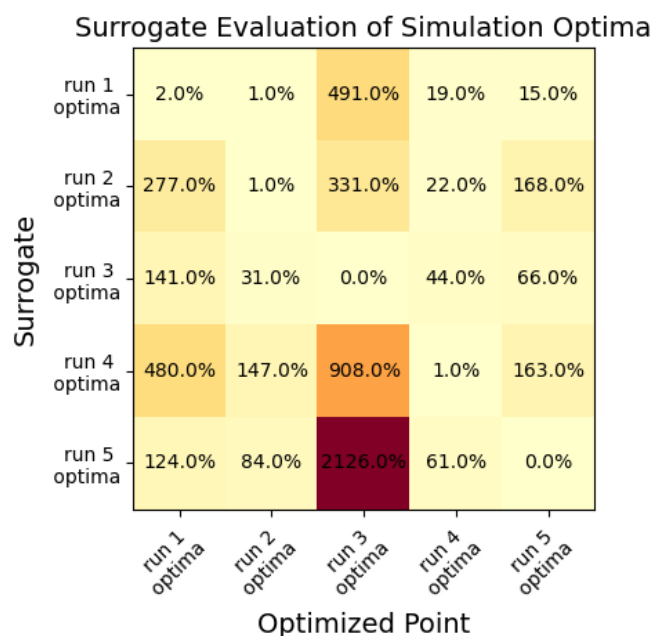


Figure 10. Cross-validation demonstrates that surrogates produced in the multi-fidelity optimization process are only locally highly predictive. Figure shows % deviation between surrogate-predicted objective functions and simulation objective functions for each of the five surrogate models and five optimization minima produced for the $N = 10$ “pure only” optimization runs.

This analysis indicates that surrogates produced as a part of a multi-fidelity optimization are usually only locally predictive, as many have very large prediction errors for objective functions at some of the other minima. The surrogate that performs best is the surrogate from run 1; which estimates the objective to within 20% of the simulation value for all cases besides the optima from run 3, which is far away from the region where other surrogates generally have samples.

We also assessed the robustness of the surrogate by performing repeated L-BFGS-B optimization on the final produced surrogates, starting from random points within the final parameter bounds box used in the optimization. This characterizes the smoothness and multimodality of the produced surrogates, as a smooth, unimodal surrogate would lead to a highly consistent local optimization, whereas a rough and

multi-modal surrogate would produce different outcomes.

For each surrogate produced in each optimization, we ran 100 L-BFGS-B optimizations from random starting points. For each of these optimizations, we calculated the standard deviation of the objective at the minimum (SD_{χ}) and the percentage of optimization within 5% of the best objective ($O_{5\%}$). The results are shown in Table 3, along with the number of simulations used to build each surrogate (K_{sim}).

Surrogate	SD_{χ}	$O_{5\%}$	K_{sim}
Optimization Run 1	0.0002	99	25
Optimization Run 2	0.017	63	20
Optimization Run 3	0.001	72	24
Optimization Run 4	0.005	94	16
Optimization Run 5	0.003	49	17

Table 3. Metrics of optimization consistency for 100 L-BFGS-B optimizations starting from random points within the bounds box for the surrogates produced in each of the 5 $N = 10$ "pure-only" optimizations. SD_{χ} is standard deviation of resulting minimized objective functions, Δ_{χ} indicates the difference between best and worst objective function, $O_{5\%}$ indicates the percentage of optimizations within 5% of the best objective for that surrogate, and N_{sim} indicates number of simulations used to build the optimization.

The results vary widely based on the surrogate, indicating that some surrogates are more robust than others. Particularly, the surrogates from optimizations 1 and 3 have the most consistent optimizations, even though optimization 3 produces a large parameter value outlier, with low standard deviations and ranges. These two surrogates also use the most simulation data and come from optimizations that more deeply explored their local optima, indicating that more sampling in the region of interest leads to a smoother, unimodal surrogate. Conversely, optimizations 2, 4 and 5 spend less time exploring the target region and have rougher surfaces, with L-BFGS-B optimizations less likely to converge. This indicates that exploring the target region in detail builds a more robust surrogate.

3.1.4 Benchmarking

We performed benchmarking on the test set described in Section 2.5 for OpenFF 1.0.0, the simulation-only optimization against the "pure only" set, and the 5 multi-fidelity optimization runs with $N = 10$ initial points. The benchmarking set is described in Section ?? We focused on the $N = 10$ runs because they generally had better objective function performance compared to the $N = 5$ runs, and they also had larger parameter changes, meaning that they would be more susceptible to overfitting. RMSE statistics for all of these force fields are plotted in Figure 11.

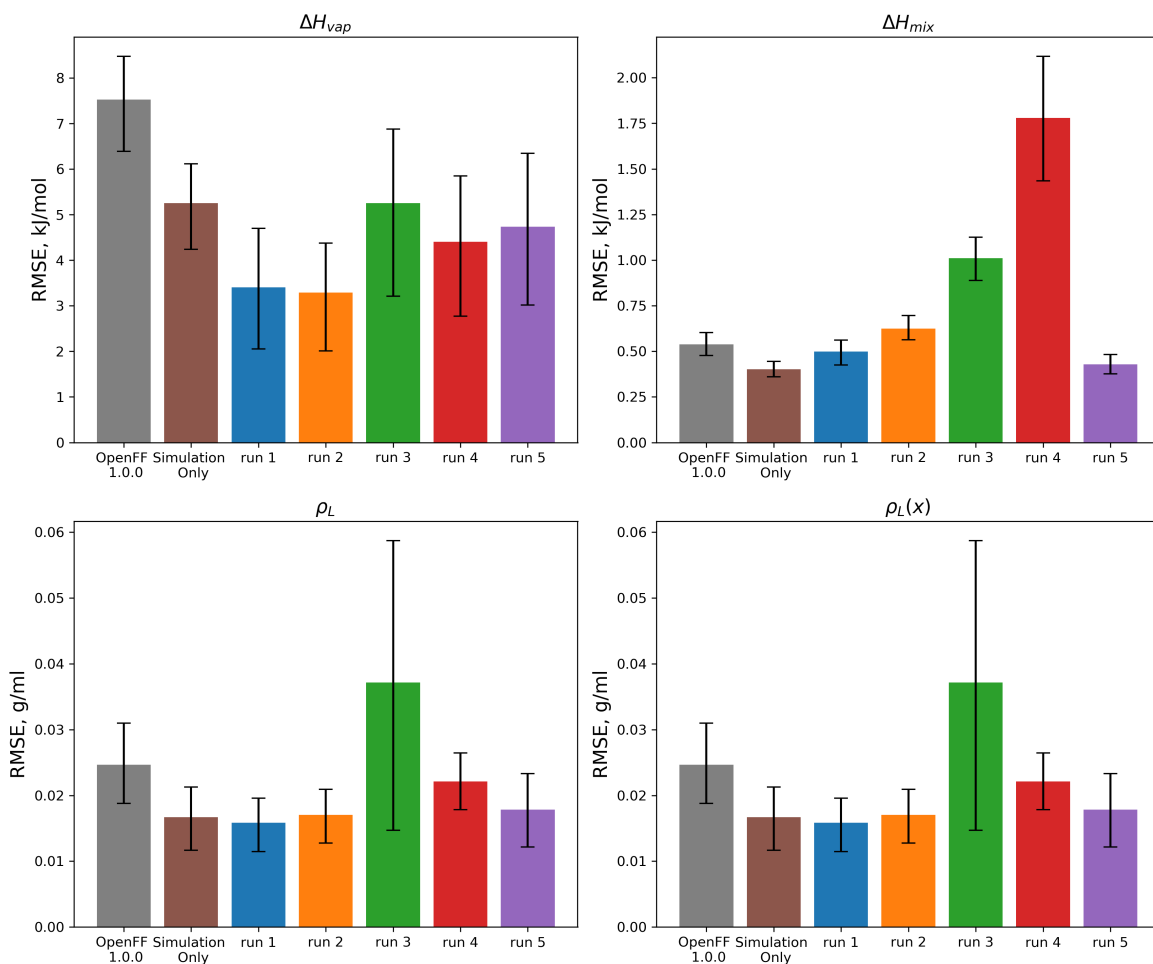


Figure 11. Test set RMSE for OpenFF 1.0.0, the simulation-only optimization, and the 5 $N = 10$ multi-fidelity runs for “pure only” targets. Benchmarking shows that some runs (such as runs 1 and 2) are transferable and outperform the simulation-only optimization, but that other runs (runs 3 and 4) have poor performance on the test set, likely due to large parameter changes in the optimization. Error bars represent 95% confidence intervals, bootstrapped over the set of molecules in the test set.

These results highlight the need to test for transferability, as run 3, which had the lowest objective function over the training set, performs worse than OpenFF 1.0.0 in three of the four physical property data types in the test set. This is caused by the very large changes in the [H-O] (hydroxyl oxygen) parameters, as this force field performs worse. Similarly, run 4 performs poorly on the test set after significant changes to the [R_{min/2}] parameter.

For optimization runs without these outlier changes to parameters, such as run 1, the results are improved, with a decrease in test set ΔH_{vap} from an initial value in OpenFF 1.0.0 of 7.52 kJ/mol (95% CI 6.42, 8.53) to a value after fitting of 3.41 kJ/mol (95% CI 1.94, 4.68), outperforming the simulation-only optimization value of 5.25 kJ/mol (95% CI 4.31, 6.16). This improvement in ΔH_{vap} following the improvement in the training set, suggests that the more aggressive multi-fidelity optimization was able to adjust parameters in a way that results in better prediction of ΔH_{vap} . For the other properties in the test set, run 1 improves over OpenFF 1.0.0 in each case, and is slightly improved over the simulation-only optimization for ρ_L and $\rho_L(x)$. The optimization does perform slightly worse than the simulation-only optimization for ΔH_{mix} , which likely reflects the lack of regularization and stronger focus on ΔH_{vap} .

We also see that some of the functional-group specific reductions in the training set RMSE translate to the test set; indeed, most of the reduction in test set RMSE comes from improved treatment of ethers and

esters, indicating that the parameter changes that led to these changes, such as significant reduction of ether ϵ , are transferable. Plots of the test set RMSE for all functional group categories are available in the Supporting Information, Section 3.1.

These results demonstrate that we can find improved parameter sets using multi-fidelity global optimization, but that care must be taken to avoid overfitting. We performed 5 optimization runs that all significantly improved the objective, but with large variations in the parameter vector solutions. There are a wide range of parameter vectors that are able to satisfy this optimization problem, but their transferability is not guaranteed. This is probably due to the training set which was chosen, as the set contains 56 target points and 12 parameters; additionally, the parameter set is “segmented” in that some parameters and targets are independent from the other parameters/targets. For example, the ether oxygen parameters are only dependent on the measurements of ρ_L and ΔH_{vap} for ethers. This leads to an optimization where overfitting is a significant concern. We could address overfitting by using a regularization scheme, as many others have done, but this many prevent us from escaping local minima in parameter space, as we hoped to do.

However, and more physically, we can also address overfitting by broadening the training set, as more physical property targets will further constrain the optimization. In addition, including mixture data in the training set helps to guard against overfitting, given that the set becomes less “segmented”, as physical properties of mixtures exercise a wider range of parameters than pure physical properties.

3.2 Mixture training set

Implementing multi-fidelity optimization with the mixture training provides us with an opportunity to test whether we can better constrain the training data without implementing regularization, since our set of mixture data is much larger, containing 195 physical property measurements. We ran an optimization with $N=20$ initial points, as well as one with $N = 10$ initial points, to determine what level of initial information was required to produce a successful optimization.

Optimization against the “mixture-only” training was successful for both the $N = 10$ and $N = 20$ runs set, achieving significant reductions in objective function and training set RMSE, as shown in Figure 12. The $N = 20$ optimization uses a total of 34 simulation evaluations to find its optimum, whereas the $N = 10$ optimization uses only 17.




Figure 12. Performance of the multi-fidelity optimization algorithm on the “mixture only” training set, for runs with $N = 10$ (blue) and $N = 20$ (orange) initial parameter vectors, indicating improved performance on ΔH_{mix} targets when compared to simulation-only optimization. Left panel shows the objective function at each iteration of the optimization. Right two panels show the RMSE for ΔH_{mix} and $\rho_L(x)$ for the two optimizations, as well as OpenFF 1.0.0 and the previous simulation-only optimization. Error bars represent 95% confidence intervals, computed with bootstrapping over the set of molecules in the training set.

In both optimizations we observe a large drop in objective, followed by incremental progress until the end of the optimization. The performance is slightly improved compared to the regularized simulation-only optimization; for the $N = 20$ run the training set ΔH_{mix} RMSE is 0.19 kJ/mol (95% CI 0.16, 0.22) vs 0.24 kJ/mol (0.21, 0.29) for the simulation-only optimization. For $\rho_L(x)$, the RMSE is 0.011 g/ml (0.01, 0.013) vs 0.013

(0.011, 0.015). Both optimizations are significantly improved when compared to OpenFF 1.0.0, with ΔH_{mix} RMSE of 0.62 kJ/mol (0.54, 0.69) and $\rho_L(x)$ RMSE of 0.23 g/ml (0.02, 0.026).

While the performance is slightly improved, the parameter changes are more significant when compared to the simulation-only optimization. The changes in parameter value from OpenFF 1.0.0 are shown in Figure 13.




Figure 13. Changes in parameter values after optimization against the “mixture only” data set, relative to OpenFF 1.0.0 (the initial values), for optimizations with both $N = 10$ and $N = 20$ initial points, as well as the simulation-only solution previously obtained. Multi-fidelity optimizations show larger parameter changes, particularly in ϵ , when compared to the simulation-only optimization over the same training set, but smaller changes in parameter values when compared to the “pure-only” multi-fidelity optimizations.

Again, the changes in parameters are larger for the multi-fidelity optimizations when compared to the simulation-only optimization, particularly for the values of ϵ . However, when compared to the multi-fidelity optimization against “pure only”, the changes are smaller and there are not significant outliers. We see some of the same parameter trends as in the “pure only” optimizations, like reduced values of ϵ for the [#1:1]-[#6X4] and [#8X2H0+0:1] atom types. A notable difference is the increased $R_{min/2}$ for the [#8:1] type, which is likely related to mixture properties better capturing the hydrogen bond donor/acceptor behavior of alcohol/ester mixtures.

3.2.1 Benchmarking

We assessed the performance of the refit force fields on the test set. Plots of RMSEs for the four types of physical property data in the test set are shown in Figure 14.




Figure 14. Test set RMSE for OpenFF 1.0.0, the simulation-only optimization, and both multi-fidelity runs, against the “mixture only” target. Benchmarking on the test set shows the transferability of the optimized parameters; notably RMSE for both ΔH_{mix} and ΔH_{vap} are significantly improved compared to OpenFF 1.0.0, despite ΔH_{vap} not being included in the training. In contrast, simulation-only optimization does not improve ΔH_{vap} . Error bars represent 95% confidence intervals, bootstrapped over the molecules in the test set.

Between the simulation-only and multi-fidelity optimizations, performance is similar on pure and mixture densities (ρ_L and $\rho_L(x)$), and not significantly improved when compared to OpenFF 1.0.0; densities are already well-predicted in OpenFF 1.0.0. For ΔH_{vap} and ΔH_{mix} , the multi-fidelity optimizations significantly improve both properties, whereas the simulation-only optimization only improved ΔH_{mix} . The $N = 20$ optimization has a ΔH_{mix} RMSE of 0.24 kJ/mol (95% CI 0.22, 0.26), similar to 0.25 (0.23, 0.27) for the simulation-only optimization. For ΔH_{vap} , the $N = 20$ optimization has an RMSE of 4.83 kJ/mol (3.75, 5.82), significantly improved over the simulation-only optimization value of 7.87 kJ/mol (6.61, 9.14), which has slightly regressed prediction of ΔH_{vap} compared to the original force field. It is notable that we are able to achieve significantly improved performance on both types of enthalpy data in the test set, indicating that parameters found with the multi-fidelity optimization process achieve better transferability than a simulation-only optimization against the same data set. Examining the test set results for the $N = 20$ run separated by functional group, as shown in Figure 15, helps to better understand how parameter changes influence force field performance.

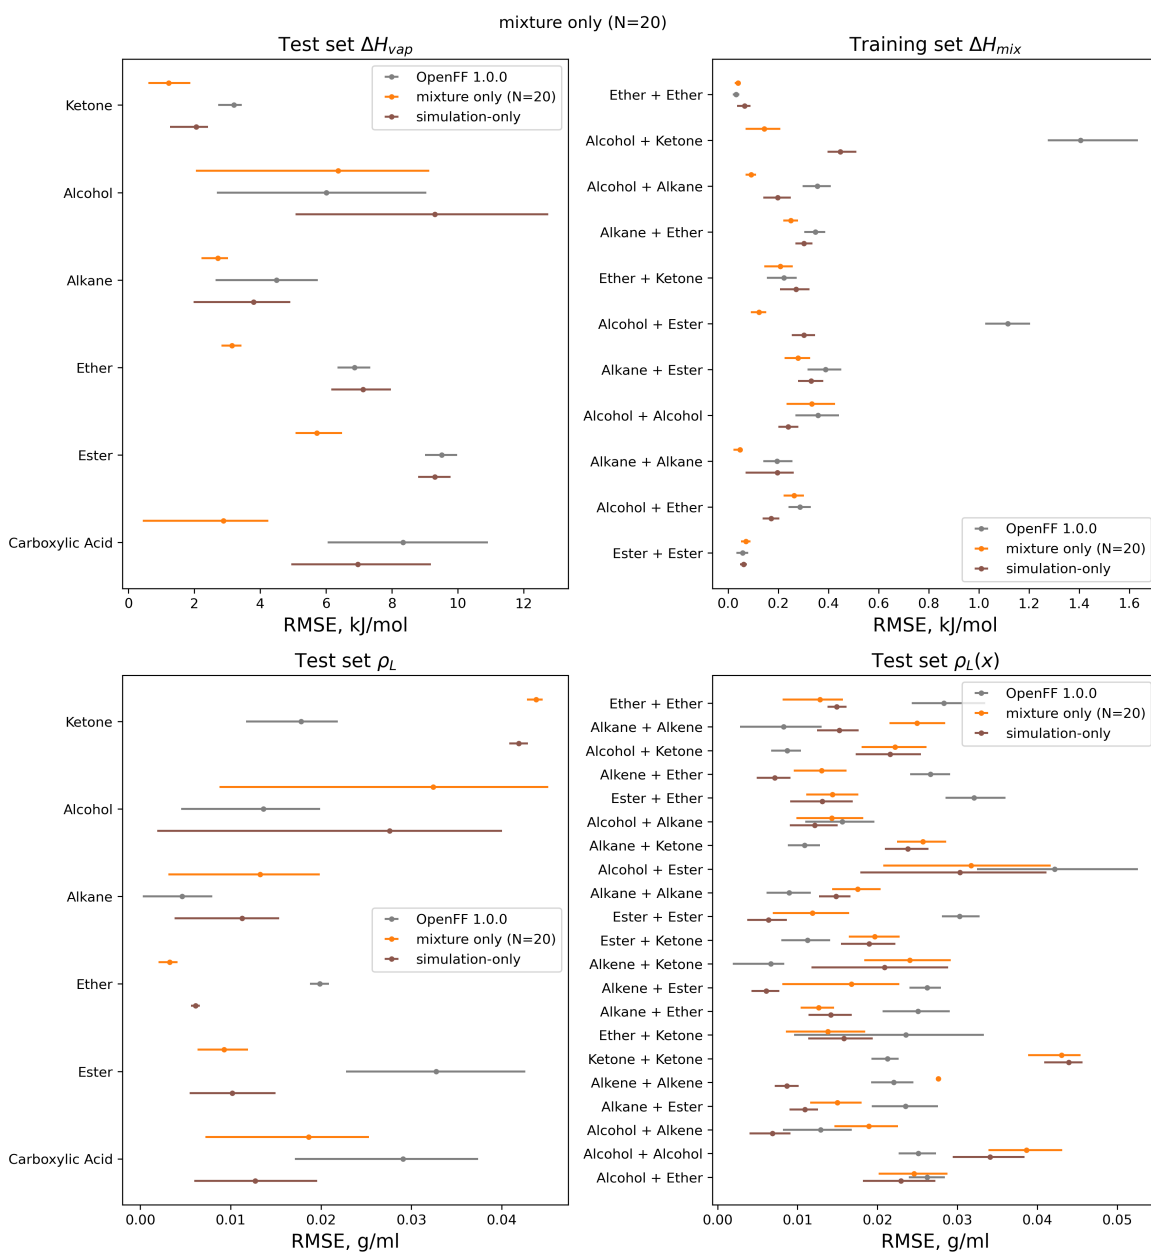


Figure 15. Test set RMSE for OpenFF 1.0.0, the simulation-only optimization, and both multi-fidelity runs, against the “mixture only” target and separated by function group or functional group pair. Benchmarking highlights important parameter changes and opportunities to tune the atom types in the model, including an apparent improvement of densities of ester mixtures at the expense of ketone mixture densities. Error bars represent 95% confidence intervals, bootstrapped over the molecules in the test set.

A notable result is that esters ρ_L and ketones perform better on ΔH_{mix} and ΔH_{vap} in this refit force field compared to their performance with the simulation-only refit. This may be due to changes in the parameters (bigger increases in $R_{min/2}$) for the [#8:1], which is exercised only by carbonyl oxygens in both the training and test sets. However, this parameter change has an interesting effect on densities; while ester densities and densities of ester-containing mixtures are largely improved, the ketone densities regressed significantly (both with multi-fidelity optimization and simulation-only optimization). This, along with parameter gradient evidence from our previous work, suggests that splitting LJ types responsible for ketones and esters may yield improved prediction of densities.

An area where the multi-fidelity optimization struggles is in the prediction of alcohols, where outside of mixtures with the improved esters/ketones, predictions are either slightly improved or degraded. While the simulation-only optimized parameters perform slightly better on alcohols, they see similar regressions in the predictions of alcohol and alcohol-mixture densities. This may be due to deficiencies in the AM1-BCC charge model for alcohols [29], leading to compensation in LJ parameters and reduced transferability.

Another type where splitting may yield improved results is the [#6:1] generic carbon type. In the training set; this type is only exercised by carbonyls, but in the test set this type is exercised by both carbonyls and alkenes. While the changes introduced in the multi-fidelity optimization significantly improve performance of carbonyl-containing molecules, binary densities of alkene mixtures are significantly degraded, indicating that the type may no longer be suitable for describing both contexts. This is sensible as carbonyls and alkenes are quite different in their chemistry.

4 Conclusions

We present a new approach for large-scale optimization force field parameters against physical property data, based on equilibrium simulations and Gaussian process surrogate modeling. Our multi-fidelity strategy uses an iterative process of global optimization over the cheap surrogate surface and validation performed at the simulation level. We demonstrate that for reasonably sized sets of physical property data, multi-fidelity optimization can find improved parameter sets while exploring more widely than traditional local optimization techniques. Particularly, we find that larger changes for parameters changes related to ethers and carbonyls yield transferable improvements, improvement test set measurements of both ΔH_{vap} and ΔH_{mix} when training against binary mixture data; training against the same dataset using only local simulation-based optimization was able to achieve comparable improvements on ΔH_{mix} , but not ΔH_{vap} . Through examination of the training and test data, we are also able to identify targets for parameter splitting.

While this strategy shows promise, challenges in implementation remain; one of the largest being the stochastic nature of the method. The improved parameters found are highly dependent on the set of initial parameter simulations used to build the surrogate model; the parameter space is rough and high-dimensional, meaning that Latin hypercube sampling struggles to find good starting sets of parameters. Building a better initial surrogate also requires more initial simulations, incurring higher computational expense. Analysis of the surrogates produced in the multi-fidelity process indicates that they are locally predictive models best suited to accelerating optimization, rather than global models accurate across the entire parameter space.

A potential route to improvement for this strategy is to incorporate Bayesian optimization[73] into the parameter search strategy in order to acquire test points more efficiently. Bayesian optimization is generally efficient at solving black-box optimization problems, which is essentially what this parameter optimization is. Starting with a smaller and more restricted set of initial parameters and allowing Bayesian optimization to acquire samples, could lead to a more efficient and reproducible optimization.

Another target area for improvement is the robustness of the surrogate building process; roughly 50% of the optimization terminate early, as issues with automatic relevance determination (ARD) cause surrogates to sacrifice accuracy to the point where they can no longer find an improved solution. While all optimization runs still led to improved parameters overall, this suggests that parameter quality could be higher with improved surrogates. Potential solutions include different choices of GP covariance kernel, or implementing alternative methods of ARD.

The production of some parameter sets with drastic changes that improve training set RMSE, but are not transferable, demonstrates that overfitting is a significant risk in parameter optimization. Typically, this risk is mitigated with regularization, penalizing solutions that stray too far from the initial solution. In this study, since we are interested in escaping local minima, we did not regularize the optimization; this led to the discovery of some significantly improved parameters (lower values of ϵ for ether oxygens, higher values of $R_{min/2}$ for carbonyl oxygens) that represented much larger changes than what regularized optimizations produced. Results from multi-fidelity optimizations on mixture properties indicates that a more complex

training target can also serve to constrain an optimization and improve transferability, while allowing parameters to vary considerably.

This surrogate-based global optimization method is able to improve force field LJ parameters by escaping local minima, leading to chemical insight and improved parameters. The success of the strategy is due to its multi-fidelity approach, using a cheaper surrogate to apply an otherwise prohibitively expensive global optimization algorithm. While already useful in its current form, the flexibility of the framework allows for significant improvement of the strategy in the future. We believe that this technique can help modelers perform better optimizations against physical property data in the future, leading to force fields which more accurately predict the behavior of molecular systems of interest.

5 Author contributions

Conceptualization: O.M., M.R.S.

Methodology: O.M.

Software: O.M.

Investigation: O.M.

Validation: O.M.

Formal Analysis: O.M.

Data Curation: O.M.

Writing - Original Draft: O.M. and M.R.S.

Writing - Review & Editing: O.M. and M.R.S.

Visualization: O.M.

Supervision: M.R.S.

Project Administration: M.R.S.

Funding Acquisition: M.R.S.

6 Data and Code Availability

To provide feedback on performance of the OpenFF force fields, we highly recommend using the issue tracker at <http://github.com/openforcefield/openforcefields>. For toolkit feedback, use <http://github.com/openforcefield/openforcefield>. Alternatively, inquiries may be e-mailed to support@openforcefield.org, though responses to e-mails sent to this address may be delayed and GitHub issues receive higher priority. For information on getting started with OpenFF, please see the documentation linked at <http://github.com/openforcefield/openforcefield>, and note the availability of several introductory examples.

7 Disclosures

MRS is an Open Science Fellow for Roivant Sciences.

Acknowledgements

We thank Simon Boothroyd for assistance with software implementation, including the OpenFF Evaluator Software; we also thank Simon for helpful conversations about the planning and scope of this project. We thank the members of the Open Force Field Initiative for their helpful feedback and discussion on the results of this research, along with the Open Force Field Scientific Advisory Board. We thank John Chodera of the Open Force Field Initiative and Sloan Kettering Institute for the use of their computational resources in this research.

Funding

Research reported in this preprint was in part supported by the National Institute of General Medical Sciences of the National Institutes of Health under award number R01GM132386. Computational resources used in this research were supported in part supported by the National Institute of General Medical Sciences

of the National Institutes of Health under award number P30CA008748, as well as the Sloan Kettering Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

We thank the Open Force Field Consortium for funding in the form of an Open Force Field Fellowship, including our industry partners as listed at the Open Force Field website (<http://openforcefield.org>). We also appreciate the Molecular Sciences Software Institute (MolSSI) for their support of the Open Force Field Initiative.

References

- [1] Halgren TA. Merck Molecular Force Field. I. Basis, Form, Scope, Parameterization, and Performance of MMFF94. *Journal of Computational Chemistry*. 1996; 17(5-6):490–519. doi: [10.1002/\(SICI\)1096-987X\(199604\)17:5/6<490::AID-JCC1>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1096-987X(199604)17:5/6<490::AID-JCC1>3.0.CO;2-P).
- [2] Vanommeslaeghe K, Hatcher E, Acharya C, Kundu S, Zhong S, Shim J, Darian E, Guvench O, Lopes P, Vorobyov I, Mackerell AD. CHARMM General Force Field: A Force Field for Drug-like Molecules Compatible with the CHARMM All-Atom Additive Biological Force Fields. *Journal of Computational Chemistry*. 2010; 31(4):671–690. doi: [10.1002/jcc.21367](https://doi.org/10.1002/jcc.21367).
- [3] Wang J, Wolf RM, Caldwell JW, Kollman PA, Case DA. Development and Testing of a General Amber Force Field. *Journal of Computational Chemistry*. 2004; 25(9):1157–1174. doi: [10.1002/jcc.20035](https://doi.org/10.1002/jcc.20035).
- [4] Riniker S. Fixed-Charge Atomistic Force Fields for Molecular Dynamics Simulations in the Condensed Phase: An Overview. *Journal of Chemical Information and Modeling*. 2018 Mar; 58(3):565–578. doi: [10.1021/acs.jcim.8b00042](https://doi.org/10.1021/acs.jcim.8b00042).
- [5] McCammon JA, Gelin BR, Karplus M. Dynamics of Folded Proteins. *Nature*. 1977 Jun; 267(5612):585–590. doi: [10.1038/267585a0](https://doi.org/10.1038/267585a0).
- [6] Zhao G, Perilla JR, Yufenyuy EL, Meng X, Chen B, Ning J, Ahn J, Gronenborn AM, Schulten K, Aiken C, Zhang P. Mature HIV-1 Capsid Structure by Cryo-Electron Microscopy and All-Atom Molecular Dynamics. *Nature*. 2013 May; 497(7451):643–646. doi: [10.1038/nature12162](https://doi.org/10.1038/nature12162).
- [7] Hollingsworth SA, Dror RO. Molecular Dynamics Simulation for All. *Neuron*. 2018 Sep; 99(6):1129–1143. doi: [10.1016/j.neuron.2018.08.011](https://doi.org/10.1016/j.neuron.2018.08.011).
- [8] Hsu WT, Ramirez DA, Sammakia T, Tan Z, Shirts MR. Identifying Signatures of Proteolytic Stability and Monomeric Propensity in O-glycosylated Insulin Using Molecular Simulation. *Journal of Computer-Aided Molecular Design*. 2022 Apr; 36(4):313–328. doi: [10.1007/s10822-022-00453-6](https://doi.org/10.1007/s10822-022-00453-6).
- [9] Vanommeslaeghe K, MacKerell AD. CHARMM Additive and Polarizable Force Fields for Biophysics and Computer-Aided Drug Design. *Biochimica et Biophysica Acta (BBA) - General Subjects*. 2015 May; 1850(5):861–871. doi: [10.1016/j.bbagen.2014.08.004](https://doi.org/10.1016/j.bbagen.2014.08.004).
- [10] Yu W, MacKerell AD. Computer-Aided Drug Design Methods. In: Sass P, editor. *Antibiotics: Methods and Protocols* Methods in Molecular Biology, New York, NY: Springer; 2017.p. 85–106. doi: [10.1007/978-1-4939-6634-9_5](https://doi.org/10.1007/978-1-4939-6634-9_5).
- [11] Warren GL, Andrews CW, Capelli AM, Clarke B, LaLonde J, Lambert MH, Lindvall M, Nevins N, Semus SF, Senger S, Tedesco G, Wall ID, Woolven JM, Peishoff CE, Head MS. A Critical Assessment of Docking Programs and Scoring Functions. *Journal of Medicinal Chemistry*. 2006 Oct; 49(20):5912–5931. doi: [10.1021/jm050362n](https://doi.org/10.1021/jm050362n).
- [12] O'Daniel PI, Peng Z, Pi H, Testero SA, Ding D, Spink E, Leemans E, Boudreau MA, Yamaguchi T, Schroeder VA, Wolter WR, Llarrull LI, Song W, Lastochkin E, Kumarasiri M, Antunes NT, Espahbodi M, Lichtenwalter K, Suckow MA, Vakulenko S, et al. Discovery of a New Class of Non- β -lactam Inhibitors of Penicillin-Binding Proteins with Gram-Positive Antibacterial Activity. *Journal of the American Chemical Society*. 2014 Mar; 136(9):3664–3672. doi: [10.1021/ja500053x](https://doi.org/10.1021/ja500053x).
- [13] Dickson CJ, Rosso L, Betz RM, Walker RC, Gould IR. GAFFlipid: A General Amber Force Field for the Accurate Molecular Dynamics Simulation of Phospholipid. *Soft Matter*. 2012 Aug; 8(37):9617–9627. doi: [10.1039/C2SM26007G](https://doi.org/10.1039/C2SM26007G).
- [14] Kurt B, Temel H. Parameterization of Boronates Using VFFDT and Paramfit for Molecular Dynamics Simulation. *Molecules*. 2020 Jan; 25(9):2196. doi: [10.3390/molecules25092196](https://doi.org/10.3390/molecules25092196).
- [15] Ghahremanpour MM, Tirado-Rives J, Jorgensen WL. Refinement of the Optimized Potentials for Liquid Simulations Force Field for Thermodynamics and Dynamics of Liquid Alkanes. *The Journal of Physical Chemistry B*. 2022 Aug; doi: [10.1021/acs.jpcc.2c03686](https://doi.org/10.1021/acs.jpcc.2c03686).

- [16] **Pastor RW**, MacKerell AD. Development of the CHARMM Force Field for Lipids. *The Journal of Physical Chemistry Letters*. 2011 Jul; 2(13):1526–1532. doi: [10.1021/jz200167q](https://doi.org/10.1021/jz200167q).
- [17] **Wang LP**, McKiernan KA, Gomes J, Beauchamp KA, Head-Gordon T, Rice JE, Swope WC, Martínez TJ, Pande VS. Building a More Predictive Protein Force Field: A Systematic and Reproducible Route to AMBER-FB15. *The Journal of Physical Chemistry B*. 2017 Apr; 121(16):4023–4039. doi: [10.1021/acs.jpbc.7b02320](https://doi.org/10.1021/acs.jpbc.7b02320).
- [18] **Qiu Y**, Smith DGA, Boothroyd S, Jang H, Hahn DF, Wagner J, Bannan CC, Gokey T, Lim VT, Stern CD, Rizzi A, Tjanaka B, Tresadern G, Lucas X, Shirts MR, Gilson MK, Chodera JD, Bayly CI, Mobley DL, Wang LP. Development and Benchmarking of Open Force Field v1.0.0—the Parsley Small-Molecule Force Field. *Journal of Chemical Theory and Computation*. 2021 Oct; 17(10):6262–6280. doi: [10.1021/acs.jctc.1c00571](https://doi.org/10.1021/acs.jctc.1c00571).
- [19] **Jorgensen WL**, Madura JD, Swenson CJ. Optimized Intermolecular Potential Functions for Liquid Hydrocarbons. *Journal of the American Chemical Society*. 1984 Oct; 106(22):6638–6646. doi: [10.1021/ja00334a030](https://doi.org/10.1021/ja00334a030).
- [20] **Horta BAC**, Merz PT, Fuchs PF, Dolenc J, Riniker S, Hünenberger PH. A GROMOS-Compatible Force Field for Small Organic Molecules in the Condensed Phase: The 2016H66 Parameter Set. *Journal of Chemical Theory and Computation*. 2016 Aug; 12(8):3825–3850. doi: [10.1021/acs.jctc.6b00187](https://doi.org/10.1021/acs.jctc.6b00187).
- [21] **Bayly CI**, Cieplak P, Cornell W, Kollman PA. A Well-Behaved Electrostatic Potential Based Method Using Charge Restraints for Deriving Atomic Charges: The RESP Model. *The Journal of Physical Chemistry*. 1993 Oct; 97(40):10269–10280. doi: [10.1021/j100142a004](https://doi.org/10.1021/j100142a004).
- [22] **Schauperl M**, Nerenberg PS, Jang H, Wang LP, Bayly CI, Mobley DL, Gilson MK. Non-Bonded Force Field Model with Advanced Restrained Electrostatic Potential Charges (RESP2). *Communications Chemistry*. 2020 Apr; 3(1):1–11. doi: [10.1038/s42004-020-0291-4](https://doi.org/10.1038/s42004-020-0291-4).
- [23] **Jakalian A**, Bush BL, Jack DB, Bayly CI. Fast, Efficient Generation of High-Quality Atomic Charges. AM1-BCC Model: I. Method. *Journal of Computational Chemistry*. 2000; 21(2):132–146. doi: [10.1002/\(SICI\)1096-987X\(20000130\)21:2<132::AID-JCC5>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1096-987X(20000130)21:2<132::AID-JCC5>3.0.CO;2-P).
- [24] **Jakalian A**, Jack DB, Bayly CI. Fast, Efficient Generation of High-Quality Atomic Charges. AM1-BCC Model: II. Parameterization and Validation. *Journal of Computational Chemistry*. 2002 Dec; 23(16):1623–1641. doi: [10.1002/jcc.10128](https://doi.org/10.1002/jcc.10128).
- [25] **Jorgensen WL**, Maxwell DS, Tirado-Rives J. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *Journal of the American Chemical Society*. 1996 Jan; 118(45):11225–11236. doi: [10.1021/ja9621760](https://doi.org/10.1021/ja9621760).
- [26] **Mohebfar M**, Johnson ER, Rowley CN. Evaluating Force-Field London Dispersion Coefficients Using the Exchange-Hole Dipole Moment Model. *Journal of Chemical Theory and Computation*. 2017 Dec; 13(12):6146–6157. doi: [10.1021/acs.jctc.7b00522](https://doi.org/10.1021/acs.jctc.7b00522).
- [27] **Boothroyd S**, Madin OC, Mobley DL, Wang LP, Chodera JD, Shirts MR. Improving Force Field Accuracy by Training against Condensed-Phase Mixture Properties. *Journal of Chemical Theory and Computation*. 2022 Jun; 18(6):3577–3592. doi: [10.1021/acs.jctc.1c01268](https://doi.org/10.1021/acs.jctc.1c01268).
- [28] **Schauperl M**, Kantonen SM, Wang LP, Gilson MK. Data-Driven Analysis of the Number of Lennard-Jones Types Needed in a Force Field. *Communications Chemistry*. 2020 Nov; 3(1):1–12. doi: [10.1038/s42004-020-00395-w](https://doi.org/10.1038/s42004-020-00395-w).
- [29] **Fennell CJ**, Wymer KL, Mobley DL. A Fixed-Charge Model for Alcohol Polarization in the Condensed Phase, and Its Role in Small Molecule Hydration. *The Journal of Physical Chemistry B*. 2014 Jun; 118(24):6438–6446. doi: [10.1021/jp411529h](https://doi.org/10.1021/jp411529h).
- [30] **Jia X**, Li P. Solvation Free Energy Calculation Using a Fixed-Charge Model: Implicit and Explicit Treatments of the Polarization Effect. *The Journal of Physical Chemistry B*. 2019 Feb; 123(5):1139–1148. doi: [10.1021/acs.jpbc.8b10479](https://doi.org/10.1021/acs.jpbc.8b10479).
- [31] **Boothroyd S**, Wang LP, Mobley DL, Chodera JD, Shirts MR. Open Force Field Evaluator: An Automated, Efficient, and Scalable Framework for the Estimation of Physical Properties from Molecular Simulation. *Journal of Chemical Theory and Computation*. 2022 Jun; 18(6):3566–3576. doi: [10.1021/acs.jctc.1c01111](https://doi.org/10.1021/acs.jctc.1c01111).
- [32] **Wang LP**, Martínez TJ, Pande VS. Building Force Fields: An Automatic, Systematic, and Reproducible Approach. *The Journal of Physical Chemistry Letters*. 2014 Jun; 5(11):1885–1891. doi: [10.1021/jz500737m](https://doi.org/10.1021/jz500737m).
- [33] **Byrd RH**, Lu P, Nocedal J, Zhu C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*. 1995 Sep; 16(5):1190–1208. doi: [10.1137/0916069](https://doi.org/10.1137/0916069).

- [34] **Boothroyd S**, Behara PK, Madin O, Hahn D, Jang H, Gapsys V, Wagner J, Horton J, Dotson D, Thompson M, Maat J, Gokey T, Wang LP, Cole D, Gilson M, Chodera J, Bayly C, Shirts M, Mobley D. Development and Benchmarking of Open Force Field 2.0.0 — the Sage Small Molecule Force Field. . 2022 Nov; doi: [10.26434/chemrxiv-2022-n2z1c](https://doi.org/10.26434/chemrxiv-2022-n2z1c).
- [35] **Alizadeh R**, Allen JK, Mistree F. Managing Computational Complexity Using Surrogate Models: A Critical Review. *Research in Engineering Design*. 2020 Jul; 31(3):275–298. doi: [10.1007/s00163-020-00336-7](https://doi.org/10.1007/s00163-020-00336-7).
- [36] **Deb K**, Roy PC, Hussein R. Surrogate Modeling Approaches for Multiobjective Optimization: Methods, Taxonomy, and Results. *Mathematical and Computational Applications*. 2021 Mar; 26(1):5. doi: [10.3390/mca26010005](https://doi.org/10.3390/mca26010005).
- [37] **Oliver MA**, Webster R. Kriging: A Method of Interpolation for Geographical Information Systems. *International Journal of Geographical Information Systems*. 1990 Jul; 4(3):313–332. doi: [10.1080/02693799008941549](https://doi.org/10.1080/02693799008941549).
- [38] **Chivatá Cárdenas I**. On the Use of Bayesian Networks as a Meta-Modelling Approach to Analyse Uncertainties in Slope Stability Analysis. *Georisk: Assessment and Management of Risk for Engineered Systems and Geohazards*. 2019 Jan; 13(1):53–65. doi: [10.1080/17499518.2018.1498524](https://doi.org/10.1080/17499518.2018.1498524).
- [39] **Dasari SK**, Cheddad A, Andersson P. Random Forest Surrogate Models to Support Design Space Exploration in Aerospace Use-Case. In: MacIntyre J, Maglogiannis I, Iliadis L, Pimenidis E, editors. *Artificial Intelligence Applications and Innovations IFIP Advances in Information and Communication Technology*, Cham: Springer International Publishing; 2019. p. 532–544. doi: [10.1007/978-3-030-19823-7_45](https://doi.org/10.1007/978-3-030-19823-7_45).
- [40] **Jiang P**, Zhou Q, Shao X. Surrogate Model-Based Engineering Design and Optimization. Springer Tracts in Mechanical Engineering, Singapore: Springer Singapore; 2020. doi: [10.1007/978-981-15-0731-1](https://doi.org/10.1007/978-981-15-0731-1).
- [41] **Nentwich C**, Engell S. Application of Surrogate Models for the Optimization and Design of Chemical Processes. In: *2016 International Joint Conference on Neural Networks (IJCNN)*; 2016. p. 1291–1296. doi: [10.1109/IJCNN.2016.7727346](https://doi.org/10.1109/IJCNN.2016.7727346).
- [42] **McBride K**, Sundmacher K. Overview of Surrogate Modeling in Chemical Process Engineering. *Chemie Ingenieur Technik*. 2019; 91(3):228–239. doi: [10.1002/cite.201800091](https://doi.org/10.1002/cite.201800091).
- [43] **Lackey BD**, Pürner M, Taracchini A, Marsat S. Surrogate Model for an Aligned-Spin Effective-One-Body Waveform Model of Binary Neutron Star Inspirals Using Gaussian Process Regression. *Physical Review D*. 2019 Jul; 100(2):024002. doi: [10.1103/PhysRevD.100.024002](https://doi.org/10.1103/PhysRevD.100.024002).
- [44] **Tapia G**, Khairallah S, Matthews M, King WE, Elwany A. Gaussian Process-Based Surrogate Modeling Framework for Process Planning in Laser Powder-Bed Fusion Additive Manufacturing of 316L Stainless Steel. *The International Journal of Advanced Manufacturing Technology*. 2018 Feb; 94(9-12):3591–3603. doi: [10.1007/s00170-017-1045-z](https://doi.org/10.1007/s00170-017-1045-z).
- [45] **Zhong J**, Wan HP, Yuan W, He M, Ren WX. Risk-Informed Sensitivity Analysis and Optimization of Seismic Mitigation Strategy Using Gaussian Process Surrogate Model. *Soil Dynamics and Earthquake Engineering*. 2020 Nov; 138:106284. doi: [10.1016/j.soildyn.2020.106284](https://doi.org/10.1016/j.soildyn.2020.106284).
- [46] **Befort BJ**, DeFever RS, Tow GM, Dowling AW, Maginn EJ. Machine Learning Directed Optimization of Classical Molecular Modeling Force Fields. *Journal of Chemical Information and Modeling*. 2021 Sep; 61(9):4400–4414. doi: [10.1021/acs.jcim.1c00448](https://doi.org/10.1021/acs.jcim.1c00448).
- [47] **Booker AJ**, Dennis JE, Frank PD, Serafini DB, Torczon V, Trosset MW. A Rigorous Framework for Optimization of Expensive Functions by Surrogates. *Structural optimization*. 1999 Feb; 17(1):1–13. doi: [10.1007/BF01197708](https://doi.org/10.1007/BF01197708).
- [48] **Storn R**, Price K. Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*. 1997 Dec; 11(4):341–359. doi: [10.1023/A:1008202821328](https://doi.org/10.1023/A:1008202821328).
- [49] **Virtanen P**, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. 2020 Mar; 17(3):261–272. doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [50] **Kramer O**. Genetic Algorithms. In: Kramer O, editor. *Genetic Algorithm Essentials Studies in Computational Intelligence*, Cham: Springer International Publishing; 2017. p. 11–19. doi: [10.1007/978-3-319-52156-5_2](https://doi.org/10.1007/978-3-319-52156-5_2).
- [51] **Balandat M**, Karrer B, Jiang DR, Daulton S, Letham B, Wilson AG, Bakshy E. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. arXiv:191006403 [cs, math, stat]. 2020 Dec; <http://arxiv.org/abs/1910.06403>.
- [52] **Neal RM**. Bayesian Learning for Neural Networks, vol. 118 of Lecture Notes in Statistics. Bickel P, Diggle P, Fienberg S, Krickeberg K, Olkin I, Wermuth N, Zeger S, editors, New York, NY: Springer New York; 1996. doi: [10.1007/978-1-4612-0745-0](https://doi.org/10.1007/978-1-4612-0745-0).

- [53] **Madin O**, Wagner J, Setiadi J, SimonBoothroyd, Thompson M, Rodríguez-Guerra J, Dotson D, Openforcefield/Openff-Evaluator: 0.3.4; 2021. doi: [10.5281/zenodo.4630739](https://doi.org/10.5281/zenodo.4630739). Zenodo.
- [54] **Boothroyd S**, Common Workflows — OpenFF Evaluator Documentation; <https://openff-evaluator.readthedocs.io/en/stable/properties/commonworkflows.html#simulation-layer>.
- [55] **Martínez L**, Andrade R, Birgin EG, Martínez JM. PACKMOL: A package for building initial configurations for molecular dynamics simulations. *Journal of Computational Chemistry*. 2009; 30(13):2157–2164. doi: [10.1002/jcc.21224](https://doi.org/10.1002/jcc.21224).
- [56] **Leimkuhler B**, Matthews C. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Applied Mathematics Research eXpress*. 2013 Jan; 2013(1):34–56. doi: [10.1093/amrx/abs010](https://doi.org/10.1093/amrx/abs010).
- [57] **Mobley D**, Bannan CC, Rizzi A, Bayly CI, Chodera JD, Lim VT, Lim NM, Beauchamp KA, Shirts MR, Gilson MK, Eastman PK. Open Force Field Consortium: Escaping Atom Types Using Direct Chemical Perception with SMIRNOFF v0.1. *bioRxiv*. 2018 Mar; p. 286542. doi: [10.1101/286542](https://doi.org/10.1101/286542).
- [58] **Bouhlef MA**, Hwang JT, Bartoli N, Lafage R, Morlier J, Martins JRRA. A Python Surrogate Modeling Framework with Derivatives. *Advances in Engineering Software*. 2019 Jul; p. 102662. doi: [10.1016/j.advengsoft.2019.03.005](https://doi.org/10.1016/j.advengsoft.2019.03.005).
- [59] **Frenkel M**, Chiroco RD, Diky V, Dong Q, Marsh KN, Dymond JH, Wakeham WA, Stein SE, Königsberger E, Goodwin ARH. XML-based IUPAC Standard for Experimental, Predicted, and Critically Evaluated Thermodynamic Property Data Storage and Capture (ThermoML) (IUPAC Recommendations 2006). *Pure and Applied Chemistry*. 2006 Jan; 78(3):541–612. doi: [10.1351/pac200678030541](https://doi.org/10.1351/pac200678030541).
- [60] **Riccardi D**, Bazyleva A, Paulechka E, Diky V, Magee JW, Kazakov AF, Townsend SA, Muzny CD, ThermoML/Data Archive. National Institute of Standards and Technology; 2021. doi: [10.18434/MDS2-2422](https://doi.org/10.18434/MDS2-2422).
- [61] **Cihlář J**, Hynek V, Svoboda V, Holub R. Heats of Vaporization of Alkyl Esters of Formic Acid. *Collection of Czechoslovak Chemical Communications*. 1976; 41(1):1–6. doi: [10.1135/cccc19760001](https://doi.org/10.1135/cccc19760001).
- [62] **Majer V**, Wagner Z, Svoboda V, Čadek V. Enthalpies of Vaporization and Cohesive Energies for a Group of Aliphatic Ethers. *The Journal of Chemical Thermodynamics*. 1980 Apr; 12(4):387–391. doi: [10.1016/0021-9614\(80\)90152-4](https://doi.org/10.1016/0021-9614(80)90152-4).
- [63] **Majer V**, Svoboda V, Hála S, Pick J. Temperature Dependence of Heats of Vaporization of Saturated Hydrocarbons C5–C8; Experimental Data and an Estimation Method. *Collection of Czechoslovak Chemical Communications*. 1979; 44(3):637–651. doi: [10.1135/cccc19790637](https://doi.org/10.1135/cccc19790637).
- [64] **Snelson A**, Skinner HA. Heats of Combustion: Sec-Propanol, 1,4-Dioxan, 1,3-Dioxan and Tetrahydropyran. *Transactions of the Faraday Society*. 1961 Jan; 57(0):2125–2131. doi: [10.1039/TF9615702125](https://doi.org/10.1039/TF9615702125).
- [65] **Svoboda V**, Uchytlová V, Majer V, Pick J. Heats of Vaporization of Alkyl Esters of Formic, Acetic and Propionic Acids. *Collection of Czechoslovak Chemical Communications*. 1980; 45(12):3233–3240. doi: [10.1135/cccc19803233](https://doi.org/10.1135/cccc19803233).
- [66] **Majer V**, Svoboda V, Uchytlová V, Finke M. Enthalpies of Vaporization of Aliphatic C5 and C6 Alcohols. *Fluid Phase Equilibria*. 1985 Jan; 20:111–118. doi: [10.1016/0378-3812\(85\)90026-3](https://doi.org/10.1016/0378-3812(85)90026-3).
- [67] **Uchytlová V**, Majer V, Svoboda V, Hynek V. Enthalpies of Vaporization and Cohesive Energies for Seven Aliphatic Ketones. *The Journal of Chemical Thermodynamics*. 1983 Sep; 15(9):853–858. doi: [10.1016/0021-9614\(83\)90091-5](https://doi.org/10.1016/0021-9614(83)90091-5).
- [68] **Byström K**, Månsson M. Enthalpies of Formation of Some Cyclic 1,3- and 1,4-Di- and Poly-Ethers: Thermochemical Strain in the –O–C–O– and –O–C–C–O– Groups. *Journal of the Chemical Society, Perkin Transactions 2*. 1982 Jan; (5):565–569. doi: [10.1039/P29820000565](https://doi.org/10.1039/P29820000565).
- [69] **Wolf G**. Thermochemische Untersuchungen an Cyclischen Ketonen. *Helvetica Chimica Acta*. 1972; 55(5):1446–1459. doi: [10.1002/hlca.19720550510](https://doi.org/10.1002/hlca.19720550510).
- [70] **Wadsö I**, Murto ML, Bergson G, Ehrenberg L, Brunvoll J, Bunnenberg E, Djerassi C, Records R. A Heat of Vaporization Calorimeter for Work at 25 Degrees C and for Small Amounts of Substances. *Acta Chemica Scandinavica*. 1966; 20:536–543. doi: [10.3891/acta.chem.scand.20-0536](https://doi.org/10.3891/acta.chem.scand.20-0536).
- [71] **Konicek J**, Wadsö I, Munch-Petersen J, Ohlson R, Shimizu A. Enthalpies of Vaporization of Organic Compounds. VII. Some Carboxylic Acids. *Acta Chemica Scandinavica*. 1970; 24:2612–2616. doi: [10.3891/acta.chem.scand.24-2612](https://doi.org/10.3891/acta.chem.scand.24-2612).
- [72] **Lipp SV**, Krasnykh EL, Verevkin SP. Vapor Pressures and Enthalpies of Vaporization of a Series of the Symmetric Linear N-Alkyl Esters of Dicarboxylic Acids. *Journal of Chemical & Engineering Data*. 2011 Apr; 56(4):800–810. doi: [10.1021/je100231g](https://doi.org/10.1021/je100231g).
- [73] **Frazier PI**. A Tutorial on Bayesian Optimization. *arXiv:180702811 [cs, math, stat]*. 2018 Jul; <http://arxiv.org/abs/1807.02811>.