

JAX-ReaxFF: A Gradient Based Framework for Extremely Fast Optimization of Reactive Force Fields

Mehmet Cagri Kaymak¹, Ali Rahnamoun², Kurt A. O’Hearn¹,
Adri C. T. van Duin³, Kenneth M. Merz, Jr.², and Hasan Metin
Aktulga¹

¹*Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA*

²*Department of Chemistry, Michigan State University, 578 S. Shaw Lane, East Lansing, MI 48824, USA*

³*Department of Mechanical Engineering, The Pennsylvania State University, University Park, State College, Pennsylvania 16802, United States*

Abstract

Molecular dynamics (MD) simulations facilitate the study of physical and chemical processes of interest. Classical models that govern the interactions between atoms lack reactivity, and quantum mechanics based methods come at steep computational costs. ReaxFF fills the gap between these two ends of the spectrum by allowing bond formation and breaking along with dynamic assignment of partial charges. To achieve realistic simulations using ReaxFF, model parameters must be optimized against the high fidelity training data that typically come from quantum mechanics. Existing parameter optimization methods for ReaxFF can be characterized as black-box techniques based on genetic algorithms or Monte-Carlo methods. Due to the stochastic behavior of these methods, the optimization process can require hundreds of thousands of error evaluations for complex parameter fitting tasks, significantly hampering the rapid development of high quality ReaxFF parameter sets. In this work, we present JAX-ReaxFF, a novel software tool that takes advantage of the modern machine learning cyber infrastructure to enable extremely fast optimization of ReaxFF parameters. By calculating the gradients of the loss function using the JAX library, we are able to utilize highly effective local optimization methods such as the limited Broyden–Fletcher–Goldfarb–Shanno (LBFGS) and Sequential Least Squares Programming (SLSQP) methods within a multi-start framework to ensure a good fit. Using the automated parallelization feature of JAX, JAX-ReaxFF can be executed on multi-core CPUs, GPUs or TPUs in a seamless way. By leveraging the gradient information and modern hardware accelerators, we are able to decrease the typical ReaxFF parame-

ter optimization time from days to mere minutes. Beyond enabling extremely fast parameter optimization, JAX-ReaxFF can also serve as a sandbox tool for ReaxFF, allowing domain scientists to start modifying the ReaxFF functional form itself for even more accurate modeling.

1 Introduction

Molecular dynamics (MD) is a widely adopted method to study physical and chemical processes at an atomistic level in a number of fields ranging from biophysics to chemistry to materials science [frenkelsmit, mdhistory1, mdhistory2]. Quantum mechanical (QM) simulations allow the geometries and energies to be predicted accurately by solving the Schrödinger’s equation. However, the computational complexity and cost of the QM based methods make them only viable for simulating small systems for rather short periods of timeframes. Molecular dynamics (MD) simulations, on the other hand, enable the study of large systems for relatively long timeframes through a number of approximations. In this approach, atomic nucleus together with its electrons is treated as a unit particle and interactions between atoms are governed by a force field (FF), which is essentially a set of parameterized mathematical equations aimed at capturing well-known atomic interactions such as bonds, valence angles, torsion, van der Waals, and Coulomb interactions. These simplifications greatly reduce the overall computational cost, but an important measure of the predictive power of force fields is their fidelity, i.e., how well they can reproduce the results of QM calculations and experimental studies. Development of high fidelity force fields relies heavily on optimization of various force field parameters based on carefully selected quantum-chemical and experimental reference data. With the help of these approximations and careful training, MD methods have proven to be successful in atomistic simulations with billions of degrees of freedom [14].

Classical MD models as implemented in highly popular MD software such as Amber [amber], LAMMPS [LAMMPS], GROMACS [gromacs4] and NAMD [namd] are based on the assumption of static chemical bonds and, in general, static charges. Therefore, they are not applicable to study phenomena where chemical reactions and charge polarization effects play a significant role. To address this gap, reactive force fields (e.g., ReaxFF [reaxff], REBO [rebo], Tersoff [tersoff1989modeling]) have been developed. These bond order potentials allow bonds to form and break throughout the simulation and they can dynamically assign partial charges to atoms using suitable charge models such as the electronegativity equalization method (EEM) [19]. The functional forms for reactive potentials are significantly more complex than their classical counterparts due to the presence of dynamic bonds and charges. For instance, ReaxFF has a formulation that contains more than 100 parameters for a simulation with 3 elements, and is about two orders of magnitude more expensive than a typical Lennard-Jones potential [lammips’benchmarks]. Consequently, training reactive force fields is an even more difficult task due to the need to capture complex phenomena such as charge distributions and reactions, and due to the

large number of parameters involved.

We focus on the ReaxFF method, which is one of the most impactful reactive force fields, if not the most impactful one [**reaxff**, **senftle**]. If there is an existing Reax force field with a similar purpose, it be sufficient to fine-tune that for the new target domain. When a new force field needs to be developed from scratch, multiple passes over the training data may be necessary, i.e., based on the quality of the resulting force field, the training data itself may need to go through revisions. In both scenarios, the training speed is crucial. As such, development of high-quality and fast optimization methods for ReaxFF has been an active research topic, first starting with the sequential one-parameter parabolic interpolation method (SOPPI) by van Duin [6], and then continuing with various global optimization methods such as genetic algorithms (GAs) [5, 13, 18], simulated annealing (SA) [11, 12], evolutionary algorithms (EAs) [29], particle swarm optimization (PSO) [8]. More recently, machine learning based search methods have been employed for this purpose [4, 10, 22].

Inspired by developments in machine learning, specifically in the field of automatic differentiation, we present a new software called JAX-ReaxFF that enables extremely fast optimization of Reax force field parameters. JAX is an auto-differentiation software by Google for high performance machine learning research [3], it can automatically differentiate native Python and NumPy functions. Leveraging this capability, JAX-ReaxFF automatically calculates the derivative of a given fitness function, which essentially measures the root mean squared deviation (RMSD) of a force field against a reference dataset, from Python-based implementation of the ReaxFF potential energy terms with respect to the set of force field parameters to be optimized. By learning the gradient information of the high dimensional parameter optimization space (which generally includes tens to over a hundred parameters), JAX-ReaxFF can employ the highly effective local optimization methods such as the Limited Memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm [31] and Sequential Least Squares Programming (SLSQP) [15] optimizer. The gradient information alone is obviously not sufficient to prevent the local optimizers from getting stuck in a local minima, but when combined with a multi-start approach, JAX-ReaxFF can greatly improve the training efficiency (measured in terms of the number of fitness function evaluations performed) and significantly reduce the time spent on optimizing ReaxFF parameters. Another important advantage of JAX is its architectural portability enabled by the XLA technology [**XLA**] used under the hood. Hence, JAX-ReaxFF can run efficiently on various architectures, including graphics processing units (GPU) and tensor processing units (TPU), by using thread parallelism and vector processing in a seamless way. As we demonstrate through extensive tests, JAX-ReaxFF can reduce the overall training time by up to three orders of magnitude compared to the state-of-the-art global optimization schemes, while achieving similar (or better) fitness scores and yielding accurate simulation results.

Beyond speeding up force field optimization, the Python based JAX-ReaxFF software provides an ideal sandbox environment for domain scientists, as they can go beyond parameter optimization and start experimenting with the func-

tional forms of the ReaxFF interactions, or add/remove interactions as desired. Since evaluating the gradient of the new functional forms with respect to atom positions gives forces, scientists are freed from the burden of coding the lengthy and bug-prone force calculation parts. Through automatic differentiation of the fitness function as explained above, parameter optimization for the new set of functional forms can be performed without any additional effort by the domain scientists. After parameter optimization, they can readily start running MD simulations to test the macro-scale properties predicted by the modified set of functional forms as a further validation test before production-scale simulations, or go back to editing the functional forms if desired results cannot be confirmed in this sandbox environment provided by JAX-ReaxFF. As such, we anticipate JAX-ReaxFF to be an indispensable tool for reactive molecular modeling and simulation.

2 Background and Related Work

Before going into the details of JAX-ReaxFF, we provide some background information on ReaxFF, and existing software for ReaxFF parameter optimizations.

2.1 ReaxFF Overview

ReaxFF divides the total potential energy of the system into various parts, including bonded and non-bonded interactions as shown in (1). The model takes atom coordinates and required force field parameters for the set of elements present in the system as input, and calculates all terms constituting the potential energy together with the corresponding forces. The derivative of each potential energy term with respect to atom coordinates gives the atomic forces that are fundamental to the MD simulation. There is a number of ReaxFF implementations with different features and architectural support such as the original Fortran Reax code [30], PuReMD [1, 2, 16], GULP [9] and LAMMPS [23].

$$E_{\text{system}} = E_{\text{bond}} + E_{\text{lon-pair}} + E_{\text{over}} + E_{\text{under}} + E_{\text{val}} + E_{\text{pen}} + E_{\text{tors}} + E_{\text{conj}} + E_{\text{Hbond}} + E_{\text{vdWaals}} + E_{\text{Coulomb}} \quad (1)$$

An important aspect of ReaxFF that separates it from classical MD models are the notions of bond orders and dynamic partial charges, which are actually not shown in (1). Bond order is a core concept in ReaxFF, it is used to determine the bond strength between pairs of atoms given their element types and distances. These pairwise bond orders are then subjected to corrections that take into account the information about all atoms surrounding each atom to obtain the predicted bonding information in a system. The corrected bond order constitutes the main input for common potential energy terms such as bond energy (E_{bond}), valence angle energy (E_{val}), and torsion angle energy (E_{tors}).




Figure 1: ReaxFF Model

However, in a dynamic bonding model, since atoms may not attain their optimal coordinations, additional terms such as lone pair ($E_{\text{lone-pair}}$), over/under-coordination (E_{over} , E_{under}), three-body penalty (E_{pen}) and four-body conjugation (E_{conj}) energies are needed. For systems with Hydrogen bonds, a special energy term (E_{Hbond}) is used. The van der Waals energy (E_{vdWaals}), which is based on the Morse potential, and the electrostatic energy term (E_{Coulomb}), which uses shielded and range-limited interactions based on dynamic charges calculated from charge models such as EEM [19], constitute the non-bonded terms in ReaxFF. Typically, bonded interactions are truncated at 5 Å, hydrogen bonds are effective up to 7.5 Å and non-bonded interactions are range limited to 10-12 Å depending on the parameterization. Figure 1 summarizes the calculations performed within a ReaxFF step.

2.2 ReaxFF Training

ReaxFF parameters are grouped by the number of atoms involved in the interaction such as single-body, two-body, three-body and four-body parameters in addition to the system-wide global parameters. Based on the distances and angles between the atoms and corresponding model parameters, bonded, 3-body, 4-body, hydrogen bond and non-bonded interaction lists are formed dynamically at each timestep. For every interaction, corresponding parameters are selected from the force field parameter set based on the interaction type and the element types of the atoms involved to calculate E_{system} . As described in detail

Type	Training Item	Target	Description
Charge	ID1 1	0.5	Charge for atom 1
Energy	ID1 - ID2/2 - ID3/3	50 kcal/mol	Energy difference
	ID1	-150 kcal/mol	
	ID3/2 - ID1/3	30 kcal/mol	
Geometry	ID1 1 2	1.25 Å	Distance between atom 1 and 2
	ID2 1 2 3	120°	Valence angle between atom 1, 2 and 3
	ID3 1 2 3 4	170°	Torsion angle between atom 1, 2, 3 and 4
Force	ID1 1	0.5 0.5 0.5	Forces on atom 1
	ID2	1.0	RMSG

Table 1: Training item types. Identifiers (Ex. ID1) are used as a reference to the molecules.

in [senftle2016reaxff], there exist parameter sets for different kinds of simulations such as combustion, aqueous systems, metals and biological systems. Even if there already is a parameter set for a simulation, it might require further tuning if the existing parameter set is not performing well for the system of interest. In some cases, the model needs to be trained from scratch which is a quite complex process. In general, as the number of elements in a parameter set increases, optimization of the force field becomes significantly harder due to the increasing number of possible element pairs, triplets and quadruplets.

ReaxFF training procedure requires three different inputs: i) *geometries*, a set of atom clusters crucial in describing the system of interest (e.g. bond stretching, angle and torsion scans, reaction transition states, crystal structures, etc.), ii) *training data*, properties of these atom clusters as calculated by appropriate QM models such as energy minimized structures, relative energies for bond/angle/torsion scans, partial charges or forces on atoms for a given structure, etc., and iii) model parameters to be optimized along with a fitness function that combines different types of training items in the following format

$$\text{Error}(m) = \sum_{i=1}^N \left(\frac{x_i - y_i}{\sigma_i} \right)^2 \quad (2)$$

where m is the model with a given set of force field parameters, x_i is the prediction by model m , y_i is the ground truth as calculated by QM, and σ_i^{-1} is the weight assigned to each training item.

Table 1 summarizes commonly used training data types and provides some examples. An energy-based training data item is formed through a linear relationship of different molecules (using their identifiers) because relative potential energies of different conformations rather than the exact potential energies are used in MD, since the chemical and physical processes are driven by the former. For structural items, molecules are expected to be energy minimized since it is essential for the force field to predict the lowest energy states correctly. For the other types, energy minimization is optional but usually preferred.

2.3 Related Work

Existing force field optimization methods for ReaxFF employ gradient-free black-box optimization methods such as Genetic Algorithms (GA) and Evolutionary Algorithms (EA). These methods enable a global search in a high dimensional parameter space, but at the expense of a high computational cost since they do not calculate the gradient of the surface of the optimization space, but rather rely on error evaluations at different points in the search space.

The earliest method Reax force field optimization is the sequential parabolic parameter interpolation method (SOPPI) [6]. The algorithm uses a well known one parameter-at-a-time approach where the consecutive single parameter searches are done until a certain convergence criteria is met. The algorithm is preferred for its simplicity but as the number of parameters increases, the number of one-parameter optimization steps needed for convergence increases drastically since only a small portion of the search space is explored at each round. Also the performance of this method is very dependent on the initial guess and the order of the parameters to be optimized in each step. Due to the drawbacks of this approach and the need for a more efficient and global optimization scheme, various global methods such as genetic algorithms (GAs) [5, 13, 18], simulated annealing (SA) [11, 12], evolutionary algorithms (EAs) [29], particle swarm optimization (PSO) [8] and machine learning based search methods [4, 10, 22, 26] have been investigated for Reax force field optimization. Due to the large number of methods reported, a thorough explanation and evaluation of these methods are out of scope of this paper, and we refer readers to studies that present a comparison of the most promising of these methods [27, 28]. At the absence of any gradient information, these methods have been proven to be successful for optimization of ReaxFF parameters. However, due to the nature of the global search algorithms, they require a high number of error evaluations and depending on the size of the problem, they could be very costly.

After the emergence of the highly optimized tools for machine learning to calculate gradients of complex functions automatically, a method called Intelligent-ReaxFF has been proposed to take advantage of these tools to train a force field for ReaxFF [10]. They use the TensorFlow library to calculate the gradients and optimize a force field. However, the method does not have the flexibility of the previously mentioned methods in terms of the training data. The force field only can be trained to match the ReaxFF potentials to the reference data. The energy differences between different geometries or charges cannot be used to train the force field which limits the usability. More general framework has been developed to provide end-to-end differentiable MD simulations [25]. Unlike this work, the main focus is to provide a unifying framework to study different potential functions including machine learning potential.

3 Proposed Method

3.1 Overview

JAX library performs auto-differentiation on native Python code, as such implementation of the ReaxFF energy expressions (see Fig. 1) in Python forms the core of JAX-ReaxFF. Once the individual energy expressions and the training error function are provided, JAX can easily calculate the gradient of the training error function with respect to the ReaxFF parameters to be optimized using its auto-diff functionality. As mentioned earlier, atomic forces can also be part of the training dataset, these can be calculated using the gradients of ReaxFF energy expressions with respect to atom positions, too.

Molecular systems used for force field training tend to have a small number of atoms compared to regular MD runs. Using a software designed for running simulations with thousands of atoms to run multiple small scale simulations introduces a lot of overhead. Optimizations such as iterative preconditioned sparse solvers to calculate the atomic charges, fast neighbor list generation algorithms and distributed computation etc. could potentially increase the run-time and results in unnecessarily complex code when the number of atoms in each geometry is small. Even though vanilla Python code tends to be slower than optimized Fortran or C code, when the auto-diff functionality, benefits of targeting small geometries and the just-in-compiled XLA support (discussed in Sect. 3.3) are considered, the advantages of JAX-ReaxFF outweighs the performance loss from not using Fortran or C.

While gradient-based optimization functionality is straight-forward to achieve using JAX as described above, there are a number of important considerations to realize an efficient (from a runtime point-of-view) and a scalable (from a memory utilization perspective) parameter optimization framework. Figure 2 gives an overview of the task-flow in JAX-ReaxFF. After the neighbor list and interactions lists are calculated based on the input geometries (Sect. 3.2, we cluster the input geometries based on the size of their interaction lists and align them properly in memory to ensure efficient single instruction multiple data (SIMD) parallelization (Sect. 3.3). After these preparation steps, the main optimization loop is executed until convergence or the maximum number of optimization steps are reached (which typically takes only tens of iterations). [HMA: I agree – some details about why energy minimization during training is important would be nice to include here, or in a subsection if it is more appropriate] [MCK: I can try but I think Ali knows more about the importance of energy minim.]. During the parameter optimization loop, some molecules might require energy minimization before the final calculation to prevent overfitting and also to tune the parameters based on more likely states of the atom clusters as the lower energy states are more likely to be observed. Hence, the main optimization loop contains a “gradient-based optimization” step followed by a “geometry optimization” step. We discuss each step in more detail in the ensuing subsections.




Figure 2: JAX-ReaxFF execution flow graph.

3.2 ReaxFF Model Implementation

In ReaxFF implementations for MD simulations, neighbor and interaction lists are created based on the atom positions and the fixed force field parameters. Due to the dynamic nature of interactions in ReaxFF, accurate and fast calculation of energy terms (especially the higher order ones such as valence angle and torsion) is critical. Differently from regular ReaxFF MD simulations, the force field is also dynamic during parameter optimization, hence adding to the challenges of developing an efficient implementation.

Pair-wise bonded interactions: We illustrate the challenges using bond order calculations as an example. As shown in Figure 1, all bonded interactions depend on the corrected bond order term. Initially, if the distance between two atoms is less than a given cutoff, typically 5 Å, the uncorrected bond orders (BO) are calculated according to Equation 3, where r_{ij} is the distance between the atom pair i - j , and p_{bo1} , p_{bo2} , p_{bo3} , p_{bo4} , p_{bo5} , p_{bo6} , r_o^σ , r_o^π , and $r_o^{\pi\pi}$ are the corresponding parameters.

$$BO'_{ij} = BO_{ij}^\sigma + BO_{ij}^\pi + BO_{ij}^{\pi\pi} = \exp \left[p_{bo1} \left(\frac{r_{ij}}{r_o^\sigma} \right)^{p_{bo2}} \right] + \exp \left[p_{bo3} \left(\frac{r_{ij}}{r_o^\pi} \right)^{p_{bo4}} \right] + \exp \left[p_{bo5} \left(\frac{r_{ij}}{r_o^{\pi\pi}} \right)^{p_{bo6}} \right] \quad (3)$$

Normally, if the uncorrected bond order is greater than a predetermined threshold, it is added to the initial bond list and subsequently bond order corrections are applied based on the neighborhood of the atoms forming the bond.

In the context of parameter optimization though, whether the pair i - j will form a bond above the given threshold also depends on the values of those parameters. Furthermore, if a given molecular structure in the training dataset requires geometry optimization (as is needed for most structural properties), atom positions change as well. JAX requires an expensive recompilation if the interaction list sizes change, we create the interaction lists once before the optimization starts and use masks to ignore the unwanted elements throughout the parameter and/or geometry optimization steps. For this purpose, for every unique element pair, the maximum possible distance where a given pair can have a valid bond order is found. If some BO related parameters are included in the optimization, values which maximize the BO term are selected from the specified parameter ranges. Then through a distance scan, the maximum possible distance is determined as the cutoff for inclusion of bond orders between that pair of elements. For geometries that require minimization, the maximum calculated distance is extended by a buffer distance to be able to accommodate potential atom position changes.

Higher Order Bonded Interactions: Similar logic is applied for other types of interactions. In a given molecule with N atoms, when there is no trimming, there will be $O(N^3)$ 3-body and $O(N^4)$ 4-body interactions. Trimming these interaction lists is required to decrease the computational and memory costs. 3-body and 4-body interaction lists are built using the corrected BO term. Another threshold is applied to the bonds forming the 3-body and 4-body interaction lists. Since the higher order bonded interactions are built using the corrected BO terms, the thresholds are also based on the previously described maximum possible BO terms. Further trimming of the lists is possible by scanning multiple distances and angles but due to the increased computational complexity, only BO term based trimming is employed.

Non-bonded Interactions: It is assumed that there is a non-bonded interaction between every atom in the system since the non-bonded interaction cutoff (which is typically 10 Å) is a lot larger than the molecular/crystal structures used for training. Therefore, non-bonded interactions form an $N * N$ matrix. If the system has periodic boundary conditions, the box dimensions are a, b and c Å and non-bonded interaction cutoff is r Å, then the size of the tensor for non-bonded interactions will become $N * N * (2 * \lceil \frac{r}{a} \rceil + 1) * (2 * \lceil \frac{r}{b} \rceil + 1) * (2 * \lceil \frac{r}{c} \rceil + 1)$. The later part is the number of periodic boxes.

Evaluation of the Potential Energy: Once interaction lists are created as described above, they stay constant throughout the optimization; the unwanted interactions in these lists are simply masked. Although this approach wastes some computational power, it avoids the expensive reneighboring, interaction list recreation and recompilation steps as force field parameters evolve. It also leads to a simplified codebase, as the interaction list generation part can be separated from the force field optimization process. The interaction list creation is

performed on the CPU using multiprocessing, regardless of whether a hardware accelerator is used for the optimization part or not.

To calculate the potential energy, a similar approach to the standalone ReaxFF code is followed with the exception of charge equilibration. The Electronegativity Equalization Method (EEM) used for distributing partial charges requires the solution of a system of linear equations (for details see [19]). For large MD simulation, the linear systems are solved using pre-conditioned iterative solvers [ohearn2020fast]. However, since the number of atoms is small for the training set structures, we use a direct LU factorization that is easier to implement and auto-differentiate.

3.3 Clustering and Alignment for SIMD Parallelization

JAX uses Accelerated Linear Algebra (XLA), a domain specific compiler for linear algebra, under the hood to achieve hardware portability. Using XLA, JAX-ReaxFF can easily run on multi-core CPUs, GPUs or TPUs without any code changes. JAX offers vectorization (*vmap*) and parallelization (*pmap*) support to take full advantage of the underlying architecture. To better utilize processing units with high level of Single Instruction Multiple Data (SIMD) parallelism, JAX offers *vmap* using which multiple small computations can be merged into batches to achieve high device utilization; *pmap* targets Multiple Instruction Multiple Data (MIMD) parallelism.

In this study, our target architecture has been GPUs as they provide significant performance advantages over multi-core CPUs and have become the mainstream hardware accelerators. However, attaining high performance on GPUs requires some important considerations. Since parameter optimization requires efficient execution of several small atomic structures as opposed to running one big MD simulation in parallel, JAX-ReaxFF leverages the *vmap* support to accelerate the energy and gradient calculations. The key for efficient vectorization in JAX-ReaxFF are the pre-calculation of interaction lists that remain static throughout optimization (as described in the previous subsection), clustering of input geometries with similar computational demands together (explained below) and alignment of the interaction lists of geometries in the same cluster (by padding as necessary) for efficient memory accesses. As mentioned before, unwanted/unnecessary interactions in these static lists are masked during the energy and gradient calculations so that they do not affect the results.

To cluster the input geometries for efficient vectorization, a modified version of the K-Means algorithm is used. The distance formula for geometry x and cluster center y with size s_y is shown below

$$\text{Dist}(x, y) = s_y(c_1 \max(n_{\{2,x\}}, n_{\{2,y\}}) + c_2 \max(n_{\{3,x\}}, n_{\{3,y\}}) + c_3 \max(n_{\{4,x\}}, n_{\{4,y\}}) + c_4 \max(n_{\{5,x\}} n_{\{1,x\}}^2, n_{\{5,y\}} n_{\{1,y\}}^2)) \quad (4)$$

where n_1, n_2, n_3, n_4 and n_5 are the number of atoms, the number of 2-body interactions, the number of 3-body interactions, the number of 4-body interactions and the number of periodic boxes within long range cutoff, respectively.

The coefficients c_1 through c_4 are indicators of the relative computational cost of their corresponding ReaxFF kernels. They can be determined empirically to accurately represent the computational costs in a given training set for a particular architecture. The cost of each cluster is determined by the computationally most expensive geometry within the cluster, as such the max value between the the current cluster center y and geometry x is chosen.

After initializing k cluster centers randomly, each geometry is assigned to these clusters based on the unique distance metric where the distance is an indicator of the change in computational load after assigning geometry x to cluster center y as shown in (4). The new center for each cluster is determined by the computationally most expensive geometry within the cluster; cluster centers determine the amount of padding needed for memory alignment of interaction lists for all geometries in their cluster. After centers are updated, a new iteration is performed where each geometry is reassigned to the closest cluster. Unlike the original K-Means algorithm, the order of geometries affects the result, therefore input geometries are shuffled after each iteration for randomization. The process continues until the cluster centers do not change anymore. Also to ensure high performance, the clustering algorithm is executed multiple times starting from different random initial cluster centers and the one where the total wasted computation (which can be determined by the total amount of padding) is minimal is chosen as the final clustering of the geometries. Although the algorithm does not guarantee optimality, empirical results are satisfactory.

Algorithm 1 Clustering Algorithm

```

1:  $C_{best} \leftarrow$  Keep track of the best so far
2: for  $r = 1, 2, \dots R$  do
3:    $C_{cur} \leftarrow$  Initialize the cluster centers by selecting a random geometry as
     the center for each cluster
4:   for  $i = 1, 2, \dots I$  do
5:      $C_{prev} \leftarrow C_{cur}$ 
6:     Shuffle  $G$ 
7:     for each  $g \in G$  do
8:       Assign  $g$  to  $c_i$  where  $\text{Dist}(g, c_i)$  is minimum
9:       Update the cluster centers
10:    end for
11:    if  $C_{cur} == C_{prev}$  then
12:      Break
13:    end if
14:  end for
15:  if  $\text{Cost}(C_{cur}) < \text{Cost}(C_{best})$  then
16:     $C_{best} \leftarrow C_{cur}$ 
17:  end if
18: end for

```

The total compilation time of JAX increases drastically with the number of

clusters because JAX unrolls the loop that iterates through the clusters during compilation. Also, if the wasted computation does not increase significantly, a smaller number of clusters is more preferable for GPUs since improving SIMD parallelism is easier within clusters. For these reasons, the number of clusters is selected automatically based on algorithm 2. Unless the computational gain from a higher number of cluster centers is not significant, smaller number of clusters is chosen.

Algorithm 2 Clustering Algorithm 2

```

1:  $k_{max} \leftarrow$  Maximum number of clusters
2:  $R \leftarrow$  Number of repetitions for the clustering algorithm
3:  $I \leftarrow$  Number of iterations for the clustering algorithm
4:  $C_{selected} \leftarrow$  Selected clustering of the geometries
5: for  $k = 1, 2, \dots, k_{max}$  do
6:    $Cost_k, C_k \leftarrow Clustering(G, k, I, R)$ 
7:   if  $|Cost_k - Cost_{k-1}|/Cost_{k-1} < tolerance$  or  $k == k_{max}$  then
8:      $C_{selected} \leftarrow C_k$ 
9:     Break
10:  end if
11: end for

```

3.4 Gradient Based Local Optimization

After the final clusters are formed, parameter optimization is performed using gradient based local optimizers with multi-start as depicted in Fig. 2. Vectorization based parallelism is employed for both energy minimization and parameter optimization steps shown in this figure.

For gradient based optimization to work, JAX traces the error function from Eq. 2 and computes the gradients of the parameters. However, since typically many geometries require geometry optimization, tracing the gradients through the optimization step is more error prone due to the complex functional form of the ReaxFF. To remedy this problem, we separate the geometry optimization from the error minimization. The error function without the geometry optimization can be thought as surrogate model since it is a fast way to approximate the true error where the geometry optimization is done as well. The approach accelerates the training and does not require tracing the gradients through the geometry optimization step.

The optimization algorithm starts from the initial geometries (Geo_{init}) and the initial force field (FF_{init}) then the force field is iteratively improved. For each iteration of the training loop shown in Algorithm-3, two different local optimizations are performed, one being local geometry optimization using the steepest descent method [**steepest descent**] and the main one being minimization of the fitness error on the energy minimized geometries by updating the prescribed force field parameters using various local optimization method such as L-BFGS-B and SLSQP. Both of these methods are classified as quasi-Newton methods

where the Hessian matrix is approximated by successive gradient calculations [15, 31]. Error minimization step uses the Geo_{min} and the optimized force field (FF_{cur}) from the last iteration and after applying the selected gradient-based algorithm, FF_{cur} gets updated with the newly trained force field. This step uses the surrogate model where the error is calculated with only the single step calculations. After that the geometry optimization step starts from Geo_{init} and yields optimized geometries (Geo_{min}) using FF_{cur} . The true error is calculated right after the geometry optimization, if there are any geometries that require it. After each iteration, the true error (E_{cur}) for FF_{cur} on the training data is calculated. If E_{cur} is lower than the lowest error so far (E_{best}), FF_{cur} is saved as the best force field (FF_{best}).

Algorithm 3 Gradient Based Local Optimization

```

1:  $FF_{cur} \leftarrow FF_{init}$ 
2: for  $iteration = 1, 2, \dots$  do
3:    $FF_{cur} \leftarrow$  Locally minimize the error through the selected gradient-based
      algorithm using  $Geo_{min}$  and starting from  $FF_{cur}$ .  $Geo_{min}$  is fixed.
4:    $Geo_{min} \leftarrow$  Geometry optimize the structures starting from the initial
      geometries  $Geo_{init}$  with the current model  $FF_{cur}$ 
5:    $E_{cur} \leftarrow$  Calculate the current error using  $Geo_{min}$  and  $FF_{cur}$ 
6:   if  $E_{cur} < E_{best}$  then
7:      $E_{best} \leftarrow E_{cur}$ 
8:      $FF_{best} \leftarrow FF_{cur}$ 
9:   end if
10:  if  $|E_{cur} - E_{prev}|/E_{prev} < 0.001$  then
11:     $FF_{cur} \leftarrow$  Add small uniform noise to  $FF_{best}$ 
12:  end if
13:   $E_{prev} \leftarrow E_{cur}$ 
14: end for

```

The error on the surrogate gets closer to the true error as the parameters converge because the changes in the parameters become minimal. One disadvantage of separating the energy minimization from the local optimization is that the fitness score for the geometry items will be ignored by the local optimization since the atom positions will not change. It introduces a discrepancy between the true error and the surrogate one. However, if the training data has multiple items related to the geometry based items as a result of potential energy surface scans (PES), the discrepancy could be minimized. As it is demonstrated in the later sections, the surrogate approach works well in practice for a variety of training tasks which have geometry based items.

4 Evaluation

In this section, we evaluate the optimization performance of JAX-ReaxFF, as well as the quality of the resulting force fields using datasets with different

characteristics.

4.1 Experimental Setup

Training Tasks For our evaluations, we identified three training tasks¹ that form a well-rounded test bench with their varying degrees of complexity. These tasks include different system types (Cobalt, a metal; Silica, an amorphous material; Disulfide, a molecular system), different types and numbers of items in the training datasets, and different number of parameters to be optimized. Also while the structures in the Cobalt and Silica training sets mostly require energy minimization, those in the Disulfide case are mostly single step energy evaluation based training items. Table 1 summarizes the specifications of the chosen training tasks. JAX-ReaxFF currently does not support training items with simulation cell optimization, as such these items are ignored. It only affects the Silica dataset which has 5 of them (out of 296 cases).

Training Data	N_{par}	N_{strc}	N_{minim}	C	G	F	P	E
Cobalt [17]	12	146	130	0	0	0	0	144
Silica [7]	67	302	221	5	26	0	6	265
Disulfide [21]	87	231	10	0	255	4401	0	219

Table 2: Datasets. N_{par} is the number of parameters to optimize, N_{strc} is the number of structures in the training dataset and N_{minim} is the number of geometries to be energy minimized. C, G, F, P and E are the number of charge based training items, geometry based items, force based items, cell parameter based items and energy based items, respectively.

Hardware Setup All the CPU experiments reported here were conducted on a computer with two Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz (2x14 cores) CPU processors, 128 GB 2133 MHz DDR4 Ram. The GPU experiments were conducted on a computer with Intel(R) Core(TM) i7-9700K CPU @ 3.60GHz, 16 GB 3000 MHz DDR4 Ram and single 1080-TI GPU card (11 GB GDDR5X memory). For the baseline methods, OGOLEM version [MCK: version] with sPuReMD version [MCK: version] backend is used. The proposed method is implemented in Python 3.7 and utilizes JAX version 0.1.76 and NumPy version 1.16.4.

Baseline Results We compare the performance and training accuracy of JAX-ReaxFF to those of methods by Shchygol et al. [27], namely the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), Monte Carlo Force Field (MCFF) optimizer and Genetic Algorithm (GA) techniques described therein.

¹The datasets could be downloaded from the following link as it is provided in the supplementary information of [27]: <https://ndownloader.figstatic.com/files/18698201>

Monte Carlo Force Field (MCFF) optimizer and Genetic Algorithm (GA) techniques described therein. MCFF optimizer utilizes simulated annealing approach to slowly modify the parameters and act based on the change in the error value. The remaining two approaches are population based and inspired by the basic principles of biological evolution. In GA, a population of candidate solutions for a given optimization problem is evolved towards better solutions. Typically, evolution happens through random mutations and cross-over between selected candidate solutions. In CMA-ES, new solutions are sampled from a multivariate normal distribution. The pairwise dependencies between the parameters are captured by the covariance matrix and as the search progresses, CMA-ES updates the covariance matrix. All three approaches use ReaxFF model as a black box and find the direction solely from the function evaluations. The settings for the black-box optimizers are given in [27]. Shchygol et al. [27] has compared these methods on different training tasks without focusing on tuning them and repeated the experiments multiple times with different starting conditions. Since they have provided an important test bench to compare different optimizers for ReaxFF, we follow the same approach to present our results.

Since the exact software and hardware from [27] are not accessible, the execution times for the baseline methods are approximated on the previously mentioned hardware. Only comparing the number of true and single step evaluations miss the execution time difference between sPuReMD backend and JAX-ReaxFF. Therefore, approximating the execution time for the baseline methods present a more detailed picture. We have calculated the time per true error evaluation using OGOLEM with sPuReMD backend and multiplied with the total number of error evaluations presented in [27]. This approximation is a lower bound for CMA-ES and MCFF since they have lower level of parallelism unlike genetic algorithm where each evaluation is independent from each other and OGOLEM uses GA for optimization.

The initial guess is an important factor which could change the results drastically. It is especially important for gradient based optimizations because these methods cannot move through the space freely as they need to follow the direction of the gradients. To show the capabilities of JAX-ReaxFF, we experimented with all three initialization methods from Shchygol et al. [27], namely random, educated and literature based initial guesses. For random initial guesses, initial values are sampled from a uniform distribution based on the given parameter ranges. To produce educated guesses, prior information from the previous related force fields is used as it is described further in [27]. For the literature based initial guesses, the force fields developed previously using the same training data are used. To give more reliable results, each initialization method is repeated ten times. For the educated and literature based initial guesses, small amount of uniform random noise is added to the parameters without violating the range restrictions. For each parameter p , the noise value is sampled from $[\frac{-1k}{10}, \frac{1k}{10}]$ where k is the length of the given range for parameter p . For the random initial guesses, uniform sampling is done ten times to produce the guesses.

4.2 Runtime and Training Evaluation

In JAX-ReaxFF, as mentioned above, two different gradient based optimization algorithms are available, L-BFGS-B and SLSQP. For both L-BFGS-B and SLSQP, the maximum number iterations is set to 100. This iteration number is for the step 3 of Algorithm 3. For L-BFGS-B, the maximum number of iterations for the line search is set to 20 and the maximum number of variable matrix corrections to approximate the Hessian matrix is set to 20. For the rest of the control parameters, the default values from the SciPy library are used. The iteration count for the main optimization loop of Algorithm 3 is set to 20 where the local error minimization and the geometry optimization steps are iteratively repeated this many times. Therefore, for all of the experiments for Jax-ReaxFF, the true error calculation with geometry optimization is done 20 times since the local error minimization only uses single step calculations.

4.2.1 Cobalt

Cobalt testcase has only energ-based training items. About 90% of these items require energy minimization, yet the training error does not fluctuate as shown in Figure 3. This shows that the surrogate error is close to the true error for this dataset. Otherwise, the error would fluctuate between iterations since the surrogate error is used for the error minimization in each iteration. For some of the random runs, SLSQP does not show any progress initially. One possible explanation is that when the initial parameters are from a non-smooth part of the optimization space that cause high gradients, the optimizer fails to escape (Figure 6b). Small noise that is added when any stall in progress is detected stimulates progress as expected.

In Table 3, we compare the convergence of JAX-ReaxFF against the the black box approaches. We observe significantly faster convergence in terms of both the number of evaluations required as well as the time taken, while obtaining similar or better training errors as measured by best and median scores. The entire optimization process takes slightly more than a minute on the GPU.




Figure 3: Convergence of the local optimizers for the Cobalt dataset

Method	Initial Guess	Best Score	Median Score	Avg. # Single Step Eval.	# True Eval.	Avg. CPU Exec. Time (min)	Avg. GPU Exec. Time (min)
L-BFGS-B	rand	1368	2334	480	20	23.5	1.2
	edu	1352	1499	418			
	lit	1366	1446	450			
SLSQP	rand	1191	2253	513	20	24.8	1.3
	edu	1168	1188	618			
	lit	1187	1189	637			
Genetic Algorithm	rand	1346	1645	-	200k	3913	-
	edu	1349	1424				
	lit	1345	1483				
CMA-ES	rand	1150	1894	-	45k	880	-
	edu	1159	1491				
	lit	1180	2320				
MCFF	rand	1422	2104	-	45k	880	-
	edu	1532	2092				
	lit	1360	1405				

Table 3: Cobalt training results.

4.2.2 Silica

The silica training set includes energy, charge and geometry matching based items. 73% of the items require energy minimization. As shown in Fig. 4, unlike the Cobalt case, the error fluctuates more between iterations, possibly because of unstable geometries as parameters are converging (note that the number of parameters to be optimized are significantly higher than the Cobalt case) and the presence of geometry matching items in the training set. Although the single point evaluation based surrogate model ignores the simulation cell optimization items, the proposed method is able to minimize the error comparable to the level of black box methods, while taking a fraction of their execution times.




Figure 4: Convergence of the local optimizers for the Silica dataset

Method	Initial Guess	Best Score	Median Score	Avg. # Single Step Eval.	# True Eval.	Avg. CPU Exec. Time (min)	Avg. GPU Exec. Time (min)
L-BFGS-B	rand	3901	5214	1865	20	25.0	1.6
	edu	4143	4467	1385			
	lit	4315	5068	1929			
SLSQP	rand	3870	4498	2962	20	31.9	2.0
	edu	3977	4540	2839			
	lit	3857	4534	2938			
Genetic Algorithm	rand	3577	3738	-	200k	1632	-
	edu	3705	3817				
	lit	3593	3721				
CMA-ES	rand	3739	4753	-	45k	367	-
	edu	3747	4122				
	lit	3793	4298				
MCFF	rand	5059	6584	-	45k	367	-
	edu	5632	7127				
	lit	4885	6126				

Table 4: Silica training results.

4.2.3 Disulfide

The disulfide training data is drastically different from the previous one since it uses force matching to fit the ReaxFF model to the data. As mentioned above, forces are calculated by taking the derivative of the potential energy expressions in JAX-ReaxFF with respect to atom positions.

$$F_x = \frac{\partial E_p}{\partial x} \quad (5)$$

$$\frac{\partial (F_x - F_t)^2}{\partial p} = \frac{\partial (\frac{\partial E_p}{\partial x} - F_t)^2}{\partial p} \quad (6)$$

where F_x is the 3-dimensional force vector for atom x , F_t is the target force vector from the training dataset and p is the model parameter to be optimized. $\frac{\partial(F_t - F_x)^2}{\partial p}$ gives the gradients for the force matching items in the objective function. However, we have observed that the final gradients for parameters from Equation 3 result in extremely high values, $\sim 10^{17}$, while the other gradients are much lower. These high gradients stop the local optimizers from doing any progress as seen in Figure 5. Since bond order parameters form the core of the dynamic bond concept in ReaxFF and therefore affect all types of bonded interactions, using relatively large parameter ranges and initial guesses being off of their ideal values are likely responsible for this behaviour. While not ideal, we excluded the bond order parameters from optimization and fixed their values to the literature ones. Among 87 parameters, 18 parameters are removed and the optimization is performed again with the remaining 69 parameters. As shown in Figure 6, this way results improve drastically, and JAX-ReaxFF can attain better scores than the baseline methods in significantly shorter time, but it should be noted that the comparison cases include all 87 parameters. This situation shows that gradient-based optimization is prone to failures for parameters with a large influence on the objective function like the bond order parameters. However, in practice training items for bond order optimization are easy to construct and their optimization can be performed independently prior to the actual optimization task. For this reason, we do not evaluate this issue to be a major limitation for JAX-ReaxFF’s practical use.




Figure 5: Convergence of the local optimizers for the Silica dataset before the modification




Figure 6: Convergence of the local optimizers for the Silica dataset after the modification

Method	Initial Guess	Best Score	Median Score	Avg. # Single Step Eval.	# True Eval.	Avg. CPU Exec. Time (min)	Avg. GPU Exec. Time (min)
L-BFGS-B*	rand	10198	10920	1660	20	9.7	0.9
	edu	10313	10631	1600			
	lit	10438	10803	1503			
SLSQP*	rand	6986	9488	1187	20	8.9	0.8
	edu	9306	9635	1234			
	lit	10304	10494	1901			
Genetic Algorithm	rand	19285	20384	-	340k	878	-
	edu	18054	20150				
	lit	18524	21206				
CMA-ES	rand	8052	11371	-	45k	116	-
	edu	8727	11105				
	lit	9284	11120				
MCFF	rand	8507	11893	-	45k	116	-
	edu	9608	13393				
	lit	10605	13625				

Table 5: Disulfide training results. *The results are from the modified version of the training.

4.3 Force Field Evaluation

Fitness scores of optimized parameters can be seen as proxies, but the quality of the resulting force field parameter sets ultimately need to be validated through actual MD simulations and comparisons against experimental and/or QM data, as we do next. MD simulations in this work are performed using Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) that is a molecular dynamics program from Sandia National Laboratories [23]. A relatively short time step of 0.5 fs was used in all the simulations. This is the recommended




Figure 7: Variations in pure Cobalt single fcc crystal cohesive energy by variations of the lattice constant.

setting for ReaxFF simulations of systems that don't include light atoms like Hydrogen. All NVT ensemble (constant number of atoms, volume and temperature) simulations were performed using Nose-Hoover thermostat to control the temperature with a temperature damping parameter of 100 fs which determines how rapidly the temperature is relaxed. All NPT ensemble (constant number of atoms, pressure and temperature) simulations were performed using Nose-Hoover thermostat to control the temperature with a temperature damping parameter of 100 fs and Nose-Hoover barostat to control the pressure with a temperature damping parameter of 1000 fs.

4.3.1 Molecular dynamics simulations of pure Cobalt structure

We investigated the crystal lattice constant correlation with cohesive energy in crystals of fcc Cobalt for validation. The lattice constant was changed from 3Å to 5Å and the associated lattice cohesive energies were recorded (Figure 7). The results of the fitted force field were compared to two previously trained ReaxFF force fields for Cobalt [17, 27] and embedded atom method (EAM) force field [24].

To validate the quality of the force field in capturing the dynamics behavior, the annealing loop was generated for a pure Cobalt crystal structure and was compared to the available force fields. A cubic simulation box of 5x5x5 ideal fcc Cobalt unit cells was generated for annealing simulations using NPT ensemble

between 1000K-3000K. After the NPT equilibration of the pure Cobalt crystal at 1000K, the system was subjected to NPT ensemble annealing between 1000K-3000K by 10 K/ps heating and cooling rate. A time step of 0.5 fs was used for the simulations. The changes in the system energy during this annealing loop is shown in Figure 8. Three ReaxFF force fields showed similar dynamic evolution behavior for the pure Cobalt structure while the EAM force field showed a different dynamic evolution (Figure 9).




Figure 8: Annealing loop of a 5x5x5 fcc Cobalt crystal including 500 atoms using fitted ReaxFF force field with heating and cooling rate of 10 K/ps.




Figure 9: Annealing loop of a 5x5x5 fcc Cobalt crystal including 500 atoms using EAM force field with heating and cooling rate of 10 K/ps.

After completion of the annealing loop, structural evaluation showed that

using the ReaxFF force fields resulted a considerable recrystallization in the pure Cobalt structure, while recrystallization was not observed when EAM force field was utilized (Figure 10).




Figure 10: Final configurations of pure Co fcc crystals after annealing loop with 1000K-3000K temperature range.

4.3.2 Molecular dynamics simulations of pure Silica structure

To evaluate the quality of the optimized force field for the Si/O parameters, the amorphous silica structure introduced in the Fogarty et al. [7] was reconstructed. The amorphous silica system included 2000 SiO_2 molecules with an initial density of 2.2 g/cm^3 (Figure 11).

The amorphous silica system was energy-minimized to eliminate initial bad contacts. The system was then annealed twice between 300K and 4000K. The first annealing loop was performed using NVT ensemble with heating and cooling rate of 37 K/ps. The second annealing loop was performed in NPT ensemble between 300K-4000K using Nose-Hoover thermostat and barostat 1.01325 bar pressure. Similar to the NVT annealing, the heating and cooling rate was 37 K/ps. Finally, the silica system was equilibrated in NPT ensemble using $T=300\text{K}$ and $P=1.01325 \text{ bar}$ for additional 200 ps as the production run. These calculations were performed using our fitted force field and two previous ReaxFF force fields introduced for such systems. The properties of the final configuration of these silica structures are compared in Table 6. The radial distributions of the final configuration of silica structure equilibrated using our fitted force field for Si-O, O-O and Si-Si are shown in Figure 12. These results show good force field quality for silica structure using JAX-ReaxFF.




Figure 11: The amorphous silica structure including 2000 SiO_2 molecules and total of 6000 atoms. Silicon atoms are shown with yellow color and Oxygen atoms are shown with red color.

Property	2010 FF [7]	2019 FF [27]	New FF
Density (g/cm^3)	2.19	2.31	2.23
Si coordination	3.99	3.94	3.97
O coordination	1.99	1.97	1.99

Table 6: Silica properties calculated using three different force fields. The experimental value reported for silica density is $2.2 g/cm^3$ [20]

4.3.3 Molecular dynamics simulations of molecules with Sulfur bonds

To test the validity of the force field containing Sulfur parameters updated using our proposed training method, we performed minimum energy structure search for single molecules with different restraints. The results of the fitted force field were compared to two previously trained ReaxFF force fields[4, 7]. The restraints are applied on S-S bond of dimethyl disulfide (DMDS), S-C bond of dimethyl thioether (DMTE), H-S-H angle of Hydrogen sulfide (H₂S) and H-S-S-H torsion angle of Hydrogen disulfide (H₂S₂). These potential energy graphs are shown in Figure 13.




Figure 12: Radial distribution function of silicon-oxygen, oxygen-oxygen, and silicon-silicon for silica structure at the end of annealing and equilibration.




Figure 13: Potential energy graphs of energy minimized molecules including Sulfur bonds with different restraints, calculated with the updated force field and previously trained force fields.

5 Conclusion

Based on the experiments, we have seen that even if the starting force field is bad, gradient based local optimizers are able to increase the fitness of the force field drastically. Clustering similar geometries together to maximize the SIMD parallelism while limiting the padding for alignment reveals high parallelism especially for single step evaluations. As it is described in Algorithm 3, by using single step energy evaluation based approximations to the error function and gradient information about the search space, we are able to decrease the convergence time significantly with the help of GPU acceleration. This allows users to generate the force fields in minutes. We have empirically showed that even though the local optimizer is not fully aware of the geometry optimization, the overall algorithm converges as the changes in parameter space becomes minimal as the algorithm progresses. Besides that, auto-diff functionality enables the study of the new functional forms for the various parts of the ReaxFF model without explicitly implementing the force calculations and optimizer since both forces and parameter gradients come for free.

References

- [1] Hasan Metin Aktulga et al. “Parallel reactive molecular dynamics: Numerical methods and algorithmic techniques”. In: *Parallel Computing* 38.4-5 (2012), pp. 245–259.
- [2] Hasan Metin Aktulga et al. “Reactive molecular dynamics: Numerical methods and algorithmic techniques”. In: *SIAM Journal on Scientific Computing* 34.1 (2012), pp. C1–C23.
- [3] James Bradbury et al. “JAX: composable transformations of Python+NumPy programs, 2018”. In: URL <http://github.com/google/jax> (2020), p. 18.
- [4] Chaitanya M Daksha et al. “Automated ReaxFF parametrization using machine learning”. In: *Computational Materials Science* 187 (), p. 110107.
- [5] Mark Dittner et al. “Efficient global optimization of reactive force-field parameters”. In: *Journal of computational chemistry* 36.20 (2015), pp. 1550–1561.
- [6] Adri CT van Duin, Jan MA Baas, and Bastiaan Van De Graaf. “Delft molecular mechanics: a new approach to hydrocarbon force fields. Inclusion of a geometry-dependent charge calculation”. In: *Journal of the Chemical Society, Faraday Transactions* 90.19 (1994), pp. 2881–2895.
- [7] Joseph C Fogarty et al. “A reactive molecular dynamics simulation of the silica-water interface”. In: *The Journal of chemical physics* 132.17 (2010), p. 174704.

- [8] David Furman et al. “Enhanced particle swarm optimization algorithm: Efficient training of reaxff reactive force fields”. In: *Journal of chemical theory and computation* 14.6 (2018), pp. 3100–3112.
- [9] Julian D Gale, Paolo Raiteri, and Adri CT van Duin. “A reactive force field for aqueous-calcium carbonate systems”. In: *Physical Chemistry Chemical Physics* 13.37 (2011), pp. 16666–16679.
- [10] Feng Guo et al. “Intelligent-ReaxFF: Evaluating the reactive force field parameters with machine learning”. In: *Computational Materials Science* 172 (2020), p. 109393.
- [11] Pierre O Hubin et al. “Parameterization of the ReaxFF reactive force field for a proline-catalyzed aldol reaction”. In: *Journal of Computational Chemistry* 37.29 (2016), pp. 2564–2572.
- [12] Eldhose Iype et al. “Parameterization of a reactive force field using a Monte Carlo algorithm”. In: *Journal of computational chemistry* 34.13 (2013), pp. 1143–1154.
- [13] Andres Jaramillo-Botero, Saber Naserifar, and William A Goddard III. “General multiobjective force field optimization framework, with application to reactive force fields for silicon carbide”. In: *Journal of Chemical Theory and Computation* 10.4 (2014), pp. 1426–1439.
- [14] Jaewoon Jung et al. “Scaling molecular dynamics beyond 100,000 processor cores for large-scale biophysical simulations”. In: *Journal of computational chemistry* 40.21 (2019), pp. 1919–1930.
- [15] Dieter Kraft et al. “A software package for sequential quadratic programming”. In: (1988).
- [16] Sudhir B Kylasa, Hasan Metin Aktulga, and Ananth Y Grama. “PuReMD-GPU: A reactive molecular dynamics simulation package for GPUs”. In: *Journal of Computational Physics* 272 (2014), pp. 343–359.
- [17] Matthew R LaBrosse, J Karl Johnson, and Adri CT van Duin. “Development of a transferable reactive force field for cobalt”. In: *The Journal of Physical Chemistry A* 114.18 (2010), pp. 5855–5861.
- [18] Henrik R Larsson, Adri CT van Duin, and Bernd Hartke. “Global optimization of parameters in the reactive force field ReaxFF for SiOH”. In: *Journal of computational chemistry* 34.25 (2013), pp. 2178–2189.
- [19] Wilfried J Mortier, Swapan K Ghosh, and S Shankar. “Electronegativity-equalization method for the calculation of atomic charges in molecules”. In: *Journal of the American Chemical Society* 108.15 (1986), pp. 4315–4320.
- [20] RL Mozzi and n BE Warren. “The structure of vitreous silica”. In: *Journal of Applied Crystallography* 2.4 (1969), pp. 164–172.
- [21] Julian Müller and Bernd Hartke. “ReaxFF reactive force field for disulfide mechanochemistry, fitted to multireference ab initio data”. In: *Journal of chemical theory and computation* 12.8 (2016), pp. 3913–3925.

- [22] Hiroya Nakata and Shandan Bai. “Development of a new parameter optimization scheme for a reactive force field based on a machine learning approach”. In: *Journal of computational chemistry* 40.23 (2019), pp. 2000–2012.
- [23] Steve Plimpton. “Fast parallel algorithms for short-range molecular dynamics”. In: *Journal of computational physics* 117.1 (1995), pp. 1–19.
- [24] GP Purja Pun and Y Mishin. “Embedded-atom potential for hcp and fcc cobalt”. In: *Physical Review B* 86.13 (2012), p. 134116.
- [25] Samuel S Schoenholz and Ekin D Cubuk. “Jax, md: End-to-end differentiable, hardware accelerated, molecular dynamics in pure python”. In: *arXiv preprint arXiv:1912.04232* (2019).
- [26] Mert Y Sengul et al. “INDEEDopt: a deep learning-based ReaxFF parameterization framework”. In: *npj Computational Materials* 7.1 (2021), pp. 1–9.
- [27] Ganna Shchygol et al. “ReaxFF Parameter Optimization with Monte-Carlo and Evolutionary Algorithms: Guidelines and Insights”. In: *Journal of Chemical Theory and Computation* 15.12 (2019), pp. 6799–6812.
- [28] Ganna Shchygol et al. “Systematic comparison of Monte Carlo Annealing and Covariance Matrix Adaptation for the optimization of ReaxFF parameters”. In: *ChemRxiv* (2018).
- [29] Tomas Trnka, Igor Tvaroska, and Jaroslav Koca. “Automated training of ReaxFF reactive force fields for Energetics of Enzymatic reactions”. In: *Journal of chemical theory and computation* 14.1 (2018), pp. 291–302.
- [30] Adri CT Van Duin et al. “ReaxFF: a reactive force field for hydrocarbons”. In: *The Journal of Physical Chemistry A* 105.41 (2001), pp. 9396–9409.
- [31] Ciyu Zhu et al. “Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization”. In: *ACM Transactions on Mathematical Software (TOMS)* 23.4 (1997), pp. 550–560.