JCTC Journal of Chemical Theory and Computation

pubs.acs.org/JCTC Article

Managing Expectations and Imbalanced Training Data in Reactive Force Field Development: An Application to Water Adsorption on Alumina

Loïc Dumortier, Céline Chizallet, Benoit Creton, Theodorus de Bruin, and Toon Verstraelen*

Cite This: https://doi.org/10.1021/acs.jctc.3c01009 Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: ReaxFF is a computationally efficient model for reactive molecular dynamics simulations that has been applied to a wide variety of chemical systems. When ReaxFF parameters are not yet available for a chemistry of interest, they must be (re)optimized, for which one defines a set of training data that the new ReaxFF parameters should reproduce. ReaxFF training sets typically contain diverse properties with different units, some of which are more abundant (by orders of magnitude) than others. To find the best parameters, one conventionally minimizes a weighted sum of squared errors over all of the data in the training set. One of the challenges in such numerical optimizations is to assign weights so that the optimized parameters represent a good compromise among all the requirements defined in the training set. This work introduces a new loss function, called Balanced Loss, and a workflow that replaces weight assignment with a more manageable procedure. The training data are divided into categories with corresponding "tolerances", *i.e.*, acceptable root-mean-square errors for the categories, which define the expectations for the optimized ReaxFF parameters. Through the Log-Sum-Exp form of Balanced Loss, the parameter optimization is also a validation of one's expectations, providing meaningful feedback that can be used to reconfigure the tolerances if needed. The new methodology is demonstrated with a nontrivial parametrization of ReaxFF for water adsorption on alumina. This results in a new force field that reproduces both the rare and frequent properties of a validation set not used for training. We also demonstrate the robustness of the new force field with a molecular dynamics simulation of water desorption from a γ-Al_2O_3 slab model.

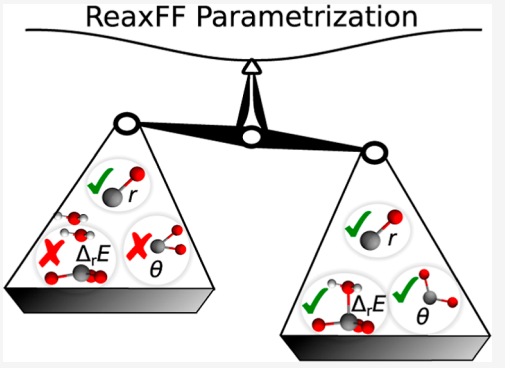

1. INTRODUCTION

Reactive force fields are widely used in molecular dynamics (MD) simulations because they combine low computational cost, close to that of molecular mechanics (MM) models, with the ability to describe chemical events, similar to more expensive quantum mechanics (QM) methods such as density functional theory (DFT). Unlike hybrid QM/MM schemes,[1] reactive force fields handle many simultaneous chemical reactions throughout the simulation cell, not just in one predefined active site. This is advantageous for the direct simulation of reaction networks[2] of complex chemical processes such as combustion,[3] pyrolysis,[4,5] chemisorption,[6,7] catalysis,[7] mechanochemistry,[8] crack propagation,[9] nucleation,[10,11] and so on. ReaxFF is one of the most established reactive force fields and is efficient enough to perform multinanosecond MD of systems with thousands of atoms using only a single high-performance compute node.[12−14] Compared to other popular reactive force fields such as Tersoff,[15] AIREBO-M,[16] or COMB3,[17] ReaxFF has been parametrized for more diverse chemical spaces.[14] More recently, machine learning potentials have also been proposed for reactive MD simulations.[18−20] All of these models share the ambition to simulate complex chemical systems at a computational cost that scales like that of MM models.

The computational efficiency of reactive force fields comes at a price. They are generally empirical models, sometimes inspired by physical principles, whose parameters must be fitted to reproduce a chemistry of interest. Such parametrization is fraught with challenges: the collection of reference datasets for training and validation, the choice of numerical optimization algorithm, the selection of parameters to optimize, the computational burden of the parameter optimization, and so on. Specifically, for ReaxFF, many optimization algorithms have been proposed and tested,[21−29] while the design of reference datasets has received much less attention. For example, ReaxFF parameters are rarely published with their training sets in a reusable form, save for a few exceptions.[8,25,27,28,30,31] However, these data are vital, as

Received: September 13, 2023
Revised: March 20, 2024
Accepted: March 22, 2024

they specify the requirements for the optimized parameters, and no model will ever outperform the data it was trained on.

A conventional ReaxFF training set consists of various target properties of relevant molecular or periodic structures, including internal coordinates, energy differences, and atomic forces. Reaction energies and barriers are obviously important for reactive force fields, but for a training set of N systems, one has at most $N - 1$ independent energy differences and many more internal coordinates and atomic forces. In the context of machine learning potentials, this imbalance is addressed by weighting data categories (typically energies and forces) inversely proportional to their prevalence,[32−34] but this practice is less established in the context of ReaxFF. Moreover, in ReaxFF, such weights are often adjusted empirically. For example, one gives more weight to an important subset of the training data in order to prioritize the performance of the trained model for that subset.[26,35,36] Conversely, the model of interest may also be inherently limited for some subsets of reference data, making it pointless to give a high weight to such subsets. These subjective motivations mean that the training set design requires expert judgment. To make this task more accessible to a broader audience, this paper introduces a new loss function and an intuitive workflow for reweighting training data called Balanced Loss. It naturally takes into account the data imbalance and the inherent strengths and weaknesses of the model being trained. A ReaxFF parametrization is used as a case study in this paper because we believe ReaxFF can greatly benefit from Balanced Loss, but the methodology is general enough to be applied to other (even nonchemical) parametrizations with similar challenges.[29,37−42] As software tools and algorithms for (re)parametrization (reactive) force fields improve,[24,26−29,43,44] we expect that more practitioners to face the challenge of data imbalance, also for machine learning potentials that are trained on increasingly large and diverse datasets.[45−47]

Alumina provides a great test case for ReaxFF parametrization because it is a versatile and widely used material in the chemical industry with a complex chemistry,[48−52] also at scales out of reach for DFT. Alumina has many known polymorphs, including γ-Al_2O_3, α-Al_2O_3, δ-Al_2O_3, and θ-Al_2O_3,[53] of which γ-Al_2O_3 is the most relevant for catalytic applications.[54,55] For example, alumina selectively adsorbs unwanted elements such as sulfur and can be used as a catalyst for the dehydration of alcohols to ethers and olefins.[56−61] However, the main application of γ-Al_2O_3 is in the automotive and petrochemical industries, where it serves as a support for other heterogeneous catalysts such as metals, metal sulfides, or metal oxides.[54,62−64] Despite their massive use in industry, the design of alumina-supported catalysts is an empirical process, partly due to the limited fundamental understanding of the materials involved. For example, the exact structure of γ-Al_2O_3 is still under discussion due to its poor crystallinity.[54,65] Also, the microscopic mechanisms at the water−alumina interface during support preparation, metal phase impregnation, shaping, and use as a catalyst remain unclear.[66]

The formation, stability, and structure of γ-Al_2O_3 are controlled by hydration and dehydration processes.[51,52,67−70] The γ phase is formed upon dehydration of boehmite at temperatures between 700 and 800 K. Once formed, the γ polymorph remains stable up to 1100 K under dry conditions.[54,55,71] γ-Al_2O_3 transitions to other polymorphs upon a further increase in temperature and/or water partial pressure. Having both Lewis acid and basic sites on the surface, alumina can react with water in several ways depending on temperature, water partial pressure (for gas/solid interfaces), and pH (for liquid/solid interfaces).[49−52,68,72] Water can adsorb without dissociation by forming a Al−O bond. One of the O−H bonds of water may then dissociate and react with a surface Al−O pair, resulting in two hydroxyl groups, called aluminols. Dissociative adsorption is reported to be more prevalent at crystal surface defects, leading to an "etching"-like degradation at these positions.[73,74] It is clear that the chemistry at the H_2O/γ-Al_2O_3 interface is highly complex and depends on the interplay of multiple microscopic mechanisms and external conditions.

Molecular simulation of the H_2O/γ-Al_2O_3 interface is a promising but ambitious method to improve our understanding of the widely used supported catalysts and to pave the way toward their rational design. DFT has often been used to model the H_2O/γ-Al_2O_3 interface.[50,52,67,69,70,75,76] Ideally, sufficiently large atomistic models are considered to avoid artificial spatial correlations, to introduce defects at low concentrations, to include both support and catalyst, and to mimic realistic water concentrations.[77−79] Because larger models also have a larger configurational space, with many local minima on the potential energy surface, their properties can no longer be simulated with static calculations, and one should resort to MD to sample all relevant configurations.[80] Linear-scaling DFT implementations[81,82] have enabled *ab initio* MD simulations of the H_2O/γ-Al_2O_3 interface,[83−85] but they are still computationally demanding compared to reactive force fields. ReaxFF has a much lower computational cost, allowing for large-scale MD simulations of alumina.[86−94] The first alumina and water ReaxFF parameters were proposed by Zhang *et al.*,[86] and these were later refined and extended by Joshi *et al.* for aluminosilicates and water,[88,90] which is particularly relevant for simulations of alumina-supported catalysts. However, as shown in the results, the state-of-the-art ReaxFF parameters by Joshi *et al.* poorly reproduce DFT reference data for water adsorption on alumina. This motivated us to demonstrate the relevance of Balanced Loss with a reparameterization of ReaxFF for H_2O/γ-Al_2O_3 interactions using DFT data from the literature.[70,76,83,95]

The rest of the paper is structured as follows. Section 2 contains the methodological details of the study: a brief overview of ReaxFF, the generation of the training and validation datasets, the parameter selection, and the optimization algorithm. The Balanced Loss function and workflow are described and motivated in detail in Section 3. Section 4 presents the results of the ReaxFF training and validation and demonstrates the suitability of the resulting force field for MD simulations. The last section formulates the main conclusions and gives an outlook on future work.

2. METHODOLOGY

2.1. ReaxFF Reactive Force Fields. ReaxFF was developed and introduced in 2001 by van Duin *et al.* for reactive MD simulations, initially of hydrocarbons,[12] and has since been regularly extended to other chemistries.[14] Like all force fields, it is a mathematical model of the interactions between atoms in a molecule or a condensed phase as a function of the Cartesian coordinates of the atomic nuclei. Unlike most classical force fields, it can describe bond breaking and formation.

The ReaxFF potential energy of an atomistic model is defined as

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

$$E_{system} = E_{bond} + E_{over} + E_{under} + E_{val} + E_{tors} + E_{vdW} + E_{charge} + E_{specific} \quad (1)$$

where E_{bond} describes the energy of an atom pair in all relevant regimes: bonded, in transition states, and dissociated. E_{over} and E_{under} are correction terms for over- and under-coordination, respectively. E_{val} is the valence angle energy and E_{tors} is the torsional angle energy between four different particles. Noncovalent interactions are modeled with E_{charge} and E_{vdW}, the charge, and the van der Waals interactions, respectively. The atomic charges are variable and account for polarization and Coulomb forces.[96,97] In addition to these commonly used energy terms, ReaxFF contains additional contributions for specific use cases, grouped into $E_{specific}$, which are not used in this work.

The covalent terms depend on bond orders (BOs), which are defined for each pair of atoms, and allow ReaxFF to describe bond breaking and formation processes in chemical reactions. The uncorrected BO of a pair of atoms consists of three terms, each corresponding to one type of covalent bond, σ, π, and $\pi\pi$

$$BO'_{ij} = BO'_{ij,\sigma} + BO'_{ij,\pi} + BO'_{ij,\pi\pi} = \exp\left(p_{bo,1}\left(\frac{r_{ij}}{r_0^{\sigma}}\right)^{p_{bo,2}}\right) + \exp\left(p_{bo,3}\left(\frac{r_{ij}}{r_0^{\pi}}\right)^{p_{bo,4}}\right) + \exp\left(p_{bo,5}\left(\frac{r_{ij}}{r_0^{\pi\pi}}\right)^{p_{bo,6}}\right) \quad (2)$$

where $\{p_{bo,n}\}_{n=1}^{6}$ represent tunable parameters that can be different for each pair of chemical elements. r_{ij} is the interatomic distance and r_0^{σ}, r_0^{π}, and $r_0^{\pi\pi}$ are the element-specific σ, π, and $\pi\pi$ equilibrium bond lengths, respectively. The expression for the uncorrected BOs in eq 2 features only a small subset of all of the adjustable ReaxFF parameters. ReaxFF has additional equations (with more parameters) to convert uncorrected to corrected BOs, which are then used in expressions for the covalent energy terms. A full description can be found in the Amsterdam Modeling Suite (AMS) documentation[98] and in the Supporting Information of ref 99.

ReaxFF was implemented in several software packages. The most established ones are the original "Standalone ReaxFF" distributed by van Duin, the commercial implementation in the AMS,[100] and the open-source version in the LAMMPS package.[101,102] In this paper, the ReaxFF implementation from AMS (release 2023.101) is used. The parameter optimization, discussed below, is implemented with ParAMS,[43,103] which is a recently developed tool in AMS for the parametrization of approximate potential energy surfaces, such as ReaxFF or density functional tight-binding (DFTB) models.[39,104] In addition, the Atomistic Simulation Environment[105] was used for processing DFT calculations in the training and validation sets. Visual Molecular Dynamics (VMD) is used for the 3D visualizations in this work.[106]

2.2. Training Set Development. The development of the training set goes through the following steps: (i) the selection of bulk (hydrated) surface and (hydrated) edge structures, (ii) periodic DFT reference calculations on these structures, and (iii) the selection of properties from these calculations as training targets.

Table 1. Overview of the Structures in the Training Set[a]

group	structure	no. of structures with additional water
bulk structures	bulk α-Al_2O_3	
	bulk boehmite	
	bulk γ-Al_2O_3	
γ-Al_2O_3 surfaces	γ-Al_2O_3(100)	4
	γ-Al_2O_3(110)	6
	γ-Al_2O_3(111)	3
boehmite surfaces	boehmite (101)	1
	boehmite (010)	0
	boehmite (100)	2
	boehmite (001)	1
γ-Al_2O_3 edges	γ-Al_2O_3(100−110)	6
small molecules	$[Al(OH)_4]^-H^+$ monomer	
	water	

[a]For surface and edge structures, the number of structures with additional water adsorbed on the surface is mentioned in the last column.

(i) **Structures**. A training set for optimizing ReaxFF parameters requires reference structures and associated training targets, such as internal coordinates or energies that ReaxFF should reproduce. The reference structures were taken from previous publications[70,76,83,95] and can be divided into five groups, as summarized in Table 1 and described in more detail below. For the sake of clarity, the relevant crystal surfaces are shown in Figure 1. Note that γ-Al_2O_3 is industrially the most relevant material, yet other forms of (hydrated) aluminum oxide were included, most notably boehmite, to increase the diversity of the training data. A complete list of structures is provided in Table S1 of the Supporting Information.

- The group of **bulk structures** contains 3D-periodic models of boehmite, γ-Al_2O_3 and α-Al_2O_3.[67]
- The group of **γ-Al_2O_3 surfaces** is based on three different slab models, cut along the (100), (110), or (111) crystal planes. In addition to the bare surfaces, structures are included with an increasing number of water molecules adsorbed on the surface.[70] All the structures were published before a distinction was made between the lateral $(110)_l$ and basal $(110)_b$ surfaces of γ-Al_2O_3, as shown in Figure 1.[52,76] The (110)γ-Al_2O_3 surfaces in the training set are in fact all lateral $(110)_l$ surfaces.
- The group of **boehmite surfaces** contains slabs with four surface orientations: (101), (010), (100), or (001). Boehmite already contains water in its bulk structure, which is preserved upon cleaving the slabs. In addition to the bare slabs, some have additional water molecules adsorbed.[83,95]
- The group of **γ-Al_2O_3 edges** comprises structures that represent the edge between surface orientations (100) and (110). In addition to the bare edge structure, six structures with an increasing number of adsorbed water molecules are included.[75]

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

- The group of **small molecules** contains two structures: a γ-alumina monomer, $[Al(OH)_4]^-H^+$, and water.[107]

(ii) **Periodic DFT calculations**. The geometries of all structures in Table 1 are optimized with DFT using periodic boundary conditions and the Perdew, Burke, and Ernzerhof (PBE) exchange−correlation functional as implemented in the Vienna *Ab initio* Simulation Package (VASP) version 5.4.[108,109] The valence interactions are described with the projected augmented wave (PAW) method.[110] An additional *a posteriori* density-dependent dispersion correction (dDsC) is applied.[111] The plane wave basis set cutoff is 600 eV, which is increased compared to the original works from which the structures were taken,[70,76,83,95] to improve the precision of the forces and facilitate the geometry optimizations. The *k*-point spacing is set to 0.5 $Å^{-1}$, and a Gaussian smearing is used with a width of 0.05 eV. The convergence criterion for the self-consistent field calculation is set to 1×10^{-5} eV. The geometry optimization is performed with the conjugate gradient algorithm and uses a convergence criterion of 0.02 eV/Å on the forces. The cell parameters were not allowed to change and were taken from previous works.[70,76,83,95]

(iii) **Property extraction**. The selection of internal coordinates, used as training targets, consists of two phases: an analysis of interatomic distances to determine appropriate cutoffs for relevant atom pairs, followed by a classification and enumeration of all relevant distances and angles.

In phase 1, histograms of all interatomic distances up to 5.0 Å were constructed per pair of chemical elements, as shown in Figure S1 in the Supporting Information. From these histograms, cutoff distances were derived to classify and enumerate all relevant atom pairs. All OH pairs with a distance below 1.2 Å are classified as covalent O−H bonds. Remaining OH distances below 2.1 Å are identified as O···H hydrogen bonds. The AlO distances below 2.8 Å are treated as Al−O bonds. AlAl distances below 4.0 Å are not directly bonded but are included because they are relevant for the local structure of alumina. No other distances were included in the training set.

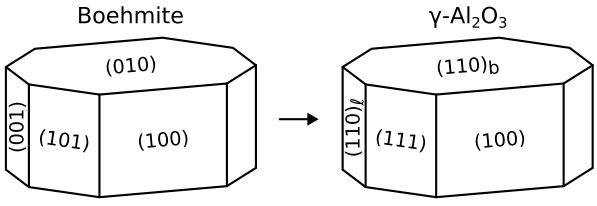

Figure 1. Overview of the crystal surfaces of boehmite and γ-Al_2O_3, with the relation between the two, as proposed in ref 52.

In phase 2, the distances defined in phase 1 are used to construct the final set of internal coordinates. In addition to distances, valence angles are constructed by combining all pairs of bonded atoms sharing one central atom. Dihedral angles are not included because most O−Al−O−X quartets, where X can be H or Al, contain nearly colinear bonds, making the dihedral angle ill-defined. Furthermore, H−O−H angles are not included because, for reasons of backward compatibility explained in the following section, the corresponding valence angle term parameters are kept fixed at the values of Joshi *et al.*[88,90] To avoid obvious redundancies in the training set, internal coordinates involving no hydrogen atoms were discarded for structures with additional water molecules, as these internal coordinates already appear in the bare surface and edge structures. The final set comprises four categories of distances (Al−O, Al−Al, O−H, and O···H) and six categories of valence angles (O−Al−O, Al−O−Al, Al−O−H, Al−O···H, H−O···H, and H···O···H). Note that the last two categories are still relevant to include, unlike H−O−H angles, because hydrogen bonding angles are sensitive to the structure of the boehmite and alumina surfaces.

Histograms of the final selection of internal coordinates can be found in Figures S2 (combined) and S3 (per material) of the Supporting Information. The internal coordinates of boehmite are somewhat similar to those of γ-Al_2O_3, with larger deviations related to hydrogen bonds. Hence, the inclusion of boehmite in the training set increases the diversity of the hydrogen bonding information. The prevalence of each class of internal coordinates per structure is given in Table S1.

Energies in a ReaxFF training set are conventionally formulated as internal energy differences between the reactant and product states. A full list of reaction energies in the training set is provided in Table S2 of the Supporting Information. These energies are grouped into five classes, denoted with three-letter codes: **BSH**, **GEH**, **GSH**, **SUR**, and **FOR**. A short summary of the included reactions and their classifications is given below.

For a given alumina surface or edge structure X, all unique pairs of adsorption states are used to construct adsorption energy training data, to avoid bias toward a particular reference state. Let m_i and $m_j > m_i$ be the number of water molecules adsorbed in states X_i and X_j, respectively, then the corresponding adsorption energy in the training set is defined as

$$\Delta_r E_{\mathrm{ads},X_jX_i} = \frac{E_{X_j} - E_{X_i}}{m_j - m_i} - E_{H_2O} \tag{3}$$

Thus, if there are N states with water adsorbed for a given structure, then there are $(N - 1)N/2$ corresponding adsorption energies in the training set. By normalizing the adsorption energies based on the number of H_2O molecules added, they all have approximately the same order of magnitude. The adsorption energies were grouped into three classes: adsorption on boehmite surfaces (**BSH**), on γ-Al_2O_3 edges (**GEH**), and on γ-Al_2O_3 surfaces (**GSH**).

While the water adsorption energies are of primary interest, other energies were also included to diversify the training set: the transformation from bulk to slab models (class **SUR**) and the formation of bulk and slab models from the alumina monomer (class **FOR**). Such "reactions" do not correspond to specific reactive events, but they do provide useful information for the covalent ReaxFF parameters.

The resulting training set contains 12,931 distances, 13,409 angles, 63 water adsorption energies, and 19 other energies, for a total of 26,422 training targets. Notably, the geometrical features, such as angles and distances, far outnumber the energy entries in the training set. The categories of internal coordinates and energies will be used in the remainder of the paper for a detailed statistical analysis and in the construction of the Balanced Loss function.

It is worth noting that our training set is considerably larger than those of previous ReaxFF parametrizations, which

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

typically contain hundreds to several thousand training targets.[8,23,28,30,31,112−114] To the best of our knowledge, the only exception is the ReaxFF model by Trnka *et al.* for enzyme chemistry, which was trained with 385,826 entries.[25] Note that their training set consists of single-point energies and forces, making it computationally less expensive, whereas ReaxFF is normally trained by optimizing geometries at each iteration of the parameter optimization.[29]

2.3. Parameter Selection. When ReaxFF parameters are calibrated, a careful selection of adjustable parameters must be made. Some parameters can be taken from previous works without refinement; others are not intended to be adjusted, such as the atomic mass and some parameters are not meaningful for the application of interest. In addition, simply optimizing all parameters would result in an intractable, high-dimensional optimization problem. The total number of ReaxFF parameters depends on the number of chemical elements in the system. The parameters are typically grouped into blocks, most of which can be repeated several times for different combinations of the chemical elements. Parameter blocks can be independent of chemical elements (41 general parameters), defined per chemical element (32 atomic parameters), per pair (16 bond and 6 off-diagonal parameters), per triplet (7 angle and 4 hydrogen-bond parameters), or per quartet of elements (7 torsion parameters).

In this work, the selection of parameters follows a top-down approach. Initially, all parameters from a literature force field are considered,[88,90] after which several selection criteria are introduced to fix parameters to their literature values, leaving only the remainder to be refined. The selection process aims for a trade-off between an acceptable dimensionality of the optimization problem and sufficient model flexibility to obtain a good fit. In the following paragraphs, we explain our selection criteria, which may be helpful for future ReaxFF calibrations.

Our starting point is the aluminosilicate force field by Joshi *et al.*[88,90] These ReaxFF parameters were calibrated to improve the description of water adsorption at acid sites, Si−O(H)−Al, in the H-ZSM-5 zeolite.[115] The literature `ffield` file contains parameters for 13 chemical elements and one dummy element, resulting in a set of 2961 parameters, many of which are irrelevant to aluminosilicates. For this work, we only consider parameter blocks that contain at least one Al element and otherwise only allow O or H. All other parameters are kept fixed, including those related to Si or those describing water.

Ideally, our reparameterization would maintain backward compatibility with the Joshi *et al.* model, changing only parameters specific to alumina and irrelevant to acid sites in zeolites. However, this severely restricts the adjustable parameters to those of atom pairs and valence angles involving at least two Al atoms, namely, the Al−Al pairs, and the Al−X−Al and Al−Al−X angles, where X can be Al, O, or H. Of this selection, only Al−O−Al angles and Al−Al pairs (≥2.5 Å) appear in the training set. As a result of this mismatch, no satisfactory reproduction of our training data was possible when imposing backward compatibility. To avoid this mismatch, we made a pragmatic selection that includes some parameters related to acid sites in zeolites but also excludes parameters that are only remotely related to our training data. This selection includes Al atom parameters, Al−X bond or pair parameters, and Al−H−O, Al−O−H, Al−O−Al, H−Al−O, and O−Al−O valence angle parameters.

We further narrow down the parameter selection using the recommendations from the ParAMS documentation.[43,103] In ParAMS, each parameter is classified with one of the following three labels: “Standard”, “Expert”, or “DoNotOptimize”. The first label indicates that the parameter is generally safe to optimize. The second label is used for parameters that should not be changed without strong motivation. Parameters with the “DoNotOptimize” label should never be touched, *e.g.*, because they contain boolean values or atomic data.[103] In this paper, only bond and angle parameters with the “Standard” label are considered for optimization.

Finally, some parameters are (de)activated for very specific reasons.

- The atomic parameter `r_0^sigma`, which is r_0^σ in eq 2, for Al is deactivated because it can be overruled by the corresponding pair parameters for Al−H and Al−O bonds.
- Because no Al−Al bonds are present, of the Al−Al parameters, only `D_e^sigma` is optimized. This introduces some freedom to tune the weak bonding interactions between pairs of Al atoms that are not directly bonded.
- π and $\pi\pi$ bond parameters are deactivated because no such bonds are present in our training data. This includes parameters with labels containing any of the following strings: `pi`, `p_bo3`, `p_bo4`, `p_bo5`, or `p_bo6`.
- The following “Expert” parameters were activated to improve the angular energy terms: `p_val3`, `p_val4`, and `p_val5`.

These selection criteria result in a subset of 36 activate parameters. For each parameter, lower and upper bounds of suitable values are determined and used to restrict the search space during the parameter optimization. For each parameter, the bounds are set equal to the corresponding range of historical values in the ReaxFF parameter database curated by Software for Chemistry & Materials B.V. (SCM).[100] Subsequently, the bounds are extended to also include a window of ±20% around the values from the Joshi *et al.* force field. Note that such choices are subjective for lack of a better alternative: there are no established defaults for the parameter bounds. The list of active parameters and their bounds can be found in Table S3 of the Supporting Information.

2.4. Optimization Settings. ReaxFF parameters are typically calibrated by minimizing a loss function with a numerical optimizer. We developed a novel loss function for this work, which will be discussed in Section 3. Here, we focus on the details of the numerical optimization algorithm.

Several optimization algorithms have been proposed to refine the ReaxFF parameters. The original method proposed by van Duin was a deterministic algorithm that optimized one parameter at a time with a parabolic extrapolation.[116] More recent algorithms are stochastic, which makes them more robust to the nontrivial structure of a standard ReaxFF loss function, such as many local minima and discontinuities.[21−25,27] These difficulties arise from the small discontinuities in the ReaxFF energy itself and the noisy sensitivity of geometry optimizations (while training) to the ReaxFF parameters.[27,44] These difficulties are still present in this work, and therefore we use a stochastic derivative-free optimizer that has proven its effectiveness, *i.e.*, the covariance matrix adaptation evolutionary strategy (CMA-ES).[27,117−119]

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

The CMA-ES settings in this work follow the best practices from the literature.[117−119] The algorithm is repeated 40 times, starting from Joshi's parameters, to reduce the risk of getting stuck in an unfavorable local minimum. These repetitions are also used to test the robustness of the new loss function proposed in Section 3. The CMA population size is set to the value recommended by Hansen, $\lfloor 4 + 3 \ln N_{par} \rfloor = 14$, where $N_{par} = 36$ is the number of activated force field parameters.[117−119] ParAMS communicates dimensionless parameters to CMA-ES by linearly transforming the original parameters so that their bounds all become $[0, 1]$. The initial CMA-ES step size in these dimensionless parameters is set to 0.2. This is sufficient to let the algorithm randomize the parameters in the first few CMA iterations, after which it starts to converge, thereby guaranteeing an initial exploration of the parameter space. Each CMA run is terminated after 1000 iterations, and the parameters with the lowest loss value are selected for further analysis.

Before evaluating the loss function in each CMA iteration, all structures in the training set are optimized with the parameters generated by CMA. The maximum number of geometry iterations is set to 500, which is much higher than the default value of 30 in ParAMS. For the training set in this work, a lower setting, such as 50, 100, or 200, produces force fields that are overfitted to this lower number of geometry steps. Each CMA run is performed on 18 cores (Intel Xeon Gold 6140), with `ParallelLevels ParameterVectors = 14` and `Jobs = 2`. This results in a slight overcommitment of the cores, which is normally not recommended, but it improves the overall efficiency in this case, which can be understood as follows. The CMA-ES algorithm synchronizes after each iteration, resulting in idle time when members of the population require different CPU times. This is generally the case for ReaxFF, since the number of required geometry steps depends strongly on the parameters. By overcommitting the cores, the idle time is reduced, resulting in a more efficient use of resources.

2.5. Validation Set.

(i) **Structures**. The validation set is taken from a dataset by Raybaud *et al.*, available on NOMAD,[52,120] containing alumina structures optimized with VASP, using the same level of theory as the training data, except that a plane-wave cutoff of 400 eV was used. This set contains 53 new γ-Al_2O_3 surface structures not used for parameter optimization with different numbers of adsorbed water molecules. The surface orientations comprise (001), (111), $(110)_l$, and $(110)_b$, as shown in Figure 1. The subscripts l and b are used to distinguish between lateral and basal surfaces, respectively, which feature different Brønsted and Lewis acid sites.[52,76] The set also includes a bulk γ-Al_2O_3 model, and an isolated water molecule was added in this work using consistent VASP settings. A complete list of structures is provided in Table S4 of the Supporting Information.

Recent developments in alumina characterization have revealed an ambiguity in the terminology used in older works. In particular, earlier spinel models considered the (100), (010), and (001) surfaces to be equivalent, but it has recently been shown from nonspinel models that this is not the case.[52] To remain consistent with published datasets and with the optimizations performed in this work, the notation remains (001) for surfaces with this orientation in the validation set and (100) for surfaces with this orientation in the training set. However, they are structurally equivalent.

(ii) **Property Extraction**. Properties are extracted using the same methodology and classification as described in Section 2.2. The resulting validation set contains 16,745 distances, 9513 angles, 101 adsorption energies, and 6 other energies, for a total of 26,365 validation targets. Figures S3, S5−S7 in the Supporting Information show the histograms of these data, whereas Table S5 lists the individual reaction energies. The final force field in this work and the original one by Joshi *et al.*[88,90] are validated by comparing these geometrical properties and energies to the VASP reference data.

3. BALANCED LOSS FUNCTION AND OPTIMIZATION WORKFLOW

ReaxFF parameters are conventionally calibrated by minimizing a loss function L, which is often a weighted sum-of-squares error (SSE) or root-mean-squared error (RMSE)

$$L_{SSE}(\mathbf{x}) = \sum_{i=1}^{N} w_i s_i(\mathbf{x}) \tag{4}$$

$$L_{RMSE}(\mathbf{x}) = \sqrt{\frac{1}{N}\sum_{i=1}^{N} w_i s_i(\mathbf{x})} \tag{5}$$

where $\{s_i\}_{i=1}^{N}$ are the squared residuals

$$s_i(\mathbf{x}) = r_i^2(\mathbf{x}) = \left(\frac{y_i - \hat{y}_i(\mathbf{x})}{\sigma_i}\right)^2 \tag{6}$$

and where the sum over i runs over all items in the training set. In every term, the property value i is calculated with a reference method (y_i) and ReaxFF ($\hat{y}_i$). Through the ReaxFF property values, the loss function depends on a vector of adjustable ReaxFF parameter vector $\mathbf{x}$. Note that CMA-ES is insensitive to the application of any monotonically increasing transformation of the loss function, so from CMA's perspective, $L_{RMSE}(\mathbf{x})$ and $L_{SSE}(\mathbf{x})$ are equivalent.

The constant σ_i is a configurable scaling factor with the same unit as the property of y_i to make the residual r_i dimensionless. In ParAMS, σ_i is only used as a reasonable order of magnitude for the corresponding y_i.[103] The weight w_i controls the importance of each training set entry in the total loss function. In principle, one can absorb w_i into σ_i or *vice versa*. The main motivation for supporting both factors in ParAMS is to cater to different user groups, some of which may prefer one over the other. This may seem surprising since textbook treatments of the least-squares method do not mention the weights w_i and only introduce σ_i as a measurement error. However, a basic assumption of the standard least-squares method does not hold here: our data have no measurement errors. Any discrepancies between the training data and ReaxFF are due to systematic errors, mainly in ReaxFF and, in principle, also in the model used to compute the training data.

At first glance, setting the weights seems straightforward: the more important an item in the training set, the higher its weight should be. However, there are different (possibly competing) motivations for adjusting the weights. The first

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

purpose of the weights is to compensate for an imbalance in the training set. For example, our training set contains many more distances and angles than energies, and the energies are also important. This imbalance can be addressed by classifying the data into categories and setting the weight to the inverse of the number of elements in each category. This strategy is common in the context of machine learning potentials, *e.g.*, when atomic forces are much more abundant than molecular energies.[32−34] A second purpose of the weights is to emphasize the importance of some residuals. For example, when an initial optimization leads to parameters for which some residuals r_i are perceived to be too large, one may increase the corresponding weights w_i and reoptimize. Assigning different weights to subsets of data is also known in the field of multiobjective optimization as the "scalarization" of multiple objectives into a single loss function.[121] Unlike scalarization methods, multiobjective evolutionary algorithms do not assume any trade-offs between categories *a priori* and instead find many Pareto-optimal solutions.[122]

While seemingly intuitive, manual weight adjustment becomes intractable when many weights need to be adjusted differently. Due to the nonlinear response of the residuals to the weights,[37] multiple combinations of weights must be tried before one reaches the residuals of interest. If the model has insufficient functional flexibility, it may be impossible to reach the desired residuals. In addition, when some residuals of interest decrease, others inevitably increase. It is difficult to predict which residuals will increase and by how much, and this can force the operator to keep adjusting the weights. In practice, this resembles a cat-and-mouse game between weights and residuals. Manually adjusting the weights also provides little insight into the optimization problem: if some residuals are large, there is no straightforward way to understand whether their weights should be increased further or whether the model is simply unable to reproduce the training data.

It is clear that the development of a training set alone is rarely sufficient to find optimal parameters. Only if the training data is completely homogeneous can one simply set all weights equal. This is typically not the case for the ReaxFF training set, which contains different types of data, such as distances, angles, and energies in this work. Assigning weights to the data is therefore an unavoidable and potentially tedious task before and during a ReaxFF parameter optimization.[26,35,36]

To simplify the tedious adjustment of weights, we introduce a new method, hereafter called "Balanced Loss". As will be demonstrated in the results, this method allows for a swift balancing of the training set, and we believe that this methodology will be equally beneficial for other optimization problems facing similar challenges. The Balanced Loss method introduces a new loss function and an intuitive workflow to balance the data and gain more insight into how well the model can reproduce subsets of the training data.

Balanced Loss requires the classification of the training data into categories. Technically, the categories are C mutually exclusive and exhaustive subsets: $S_c\ \forall\ c \in \{1...C\}$. They should be defined so that residuals within a category respond in roughly the same way to a change in the model parameters. For example, one might expect that all the O−H bond lengths in a training set, while not exactly the same, do respond similarly to changes in the ReaxFF parameters. With this partitioning, the Balanced Loss function is defined as

$$L_{\mathrm{BL}} = \tau f^{-1}\left(\sum_{c=1}^{C} f\left(\frac{R_c}{\tau}\right)\right) \tag{7}$$

where R_c is the RMSE on the entries in category c

$$R_c = \sqrt{\frac{1}{|S_c|}\sum_{i\in S_c} s_i} \tag{8}$$

L_{BL} is dimensionless by construction. The function f and its inverse must be monotonically increasing functions, and by default $f(x) = \exp(x)$ is used, which will be denoted as the Log-Sum-Exp (LSE) form, referring to the mathematical operations in eq 7. To illustrate the benefits of the LSE form, all parametrizations will be repeated with two other forms of f: $f(x) = x^2$, denoted as root-sum-square (RSS), and $f(x) = x$, denoted as identity-sum-identity (ISI). Note that the RSS form makes Balanced Loss formally equivalent with a standard loss function in eq 5, with $w_i = \frac{N}{S_c}$, where c is the category to which training set item i belongs.

By default, L_{BL} is thus a LSE function, a well-established smooth approximation of the maximum over multiple inputs. It is popular in the machine learning context[123] and it has been used for the scalarization of multiobjective problems.[124] Here, the inputs to LSE are all R_c values. The hyperparameter τ, sometimes called the effective temperature, controls the smoothness of the approximation of the maximum. In the "cold" limit $\tau \to 0$, LSE loses its smoothness and reduces to the maximum over all R_c. The parameter τ appears in two places, such that $L_{\mathrm{BL}} = L_{\mathrm{RMSE}}$ in the trivial case of one category and $w_i = 1\ \forall\ i$. Throughout this paper, we have used $\tau = 1$.

Unlike L_{SSE} and L_{RMSE}, user-defined weights w_i are missing from L_{BL}, which implies that σ_i must play a slightly different role. We propose to set each σ_i to the desired accuracy of the corresponding entry y_i in the training set, which resembles its meaning in conventional least-squares methods. To make the distinction with σ_i in other contexts, we call them "tolerances" in the context of Balanced Loss because they represent the magnitudes of residuals one is willing to tolerate. At this stage, simply defining tolerances may seem like wishful thinking, but it will become clear later that Balanced Loss helps find a consistent set of ReaxFF parameters and tolerances. Our definition of σ_i (as the desired accuracy) also facilitates the interpretation of R_c: it expresses, in the RMS sense, the average ratio between actual and desired accuracy. In the ideal case, after optimizing the parameters, one obtains $R_c = 1\ \forall\ c$.

Given the interpretation of R_c, the LSE form of the Balanced Loss is easily motivated. If one of the R_c values is much higher than all others, one finds $L_{\mathrm{BL}} \approx R_c$, *i.e.*, category c dominates the loss function. If the optimization algorithm explores a region of the parameter space where category c dominates, then it will focus only on reducing R_c, with other categories acting at best as a form of regularization. This is a desirable feature, since category c is then the worst reproduced subset of the training data and therefore deserves the optimizer's full attention.

One may also understand the effect of LSE by comparing the gradients of L_{BL} and L_{RMSE} with respect to the ReaxFF parameters

$$\frac{\partial L_{\mathrm{BL}}}{\partial x_k} = \sum_{i=1}^{N} \frac{\partial L_{\mathrm{BL}}}{\partial s_i}\frac{\partial s_i}{\partial x_k} \tag{9}$$

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

$$\frac{\partial L_{\mathrm{RMSE}}}{\partial x_k} = \sum_{i=1}^{N} \frac{\partial L_{\mathrm{RMSE}}}{\partial s_i} \frac{\partial s_i}{\partial x_k} \tag{10}$$

Both loss gradients are linear combinations of the gradients of squared residuals, $\frac{\partial s_i}{\partial x_k}$, but they combine them with different "weights"

$$\frac{\partial L_{\mathrm{BL}}}{\partial s_i} = \frac{\exp\left(\frac{R_d}{\tau}\right)}{\sum_{c=1}^{C} \exp\left(\frac{R_c}{\tau}\right)} \frac{\tau}{2|S_d|R_d} \qquad \text{with } i \in S_d \tag{11}$$

$$\frac{\partial L_{\mathrm{RMSE}}}{\partial s_i} = \frac{w_i}{2NL_{\mathrm{RMSE}}} \tag{12}$$

In the case of L_{RMSE}, the weight $\frac{\partial L_{\mathrm{RMSE}}}{\partial s_i}$ is simply proportional to the user-defined weight w_i. For Balanced Loss, however, the weight $\frac{\partial L_{\mathrm{BL}}}{\partial s_i}$ contains a new and crucial factor

$$P_d = \frac{\exp \frac{R_d}{\tau}}{\sum_{c=1}^{C} \exp \frac{R_c}{\tau}} \tag{13}$$

with

$$\sum_{d=1}^{C} P_d = 1 \tag{14}$$

This factor is known as SoftMax, a continuous generalization of the ArgMax function, used to identify the position of a maximum in an ordered list.[125] This shows how Balanced Loss borrows a strategy from reinforcement learning, known as the Gradient Bandit Algorithm. At each iteration in the optimization, the most violated subset of the training data determines the action,[126] in this context, action being the direction in which the parameters must evolve. The analogy between P_d (in L_{BL}) and user-defined weights w_i (in L_{RMSE}) also suggests another interpretation. Instead of a human operator tuning the weights w_i, as in the cat-and-mouse metaphor introduced above, Balanced Loss adjusts the weights algorithmically within a single optimization run.

So far, we have assumed that one simply sets the tolerance σ_i to the desired accuracy of the corresponding y_i. However, such a choice may be subjective and incompatible with the capabilities of the ReaxFF model to be trained. In practice, we recommend such "naive" tolerances σ_i as a first guess. Parameter optimization can then be used to test these expectations. To do so, we recommend the following workflow:

W1 First, all the elements of a conventional parameter optimization are gathered: (i) the model, (ii) the training data, (iii) a selection of parameters to optimize and their bounds, (iv) an initial guess of the parameters, and (v) an optimization algorithm. In this paper, all these elements are described in Section 2.

W2 Then the additional elements needed for Balanced Loss are defined: the categories of training data and an initial configuration of the tolerances σ_i. Such choices are domain-specific, but a few general recommendations can be given, in addition to the ones discussed above. It is convenient to have data with consistent units within one category, c and to assign the same tolerance to all its members, for which the symbol σ_c will be used below. Furthermore, it is useful to introduce categories for data that deserve special attention, *e.g.*, with key properties for the intended application of the force field, or with properties that are harder to reproduce than others. By placing these data in separate categories, their RMSEs are easily monitored, and large errors within these categories will be prioritized during the optimization.

For ReaxFF, one can introduce different categories for distances, angles, and energies. In this paper, the categories are more fine-grained: all internal coordinates are classified by the chemical elements and bond types involved. In fact, we categorize all training data as they were introduced in Section 2.2: the 4 bond categories are Al−O, Al−Al, O−H, and O···H, the 6 angle categories are O−Al−O, Al−O−Al, Al−O−H, Al−O···H, H−O···H, and H···O···H, and the energy categories are **BSH**, **GEH**, **GSH**, **SUR**, and **FOR**. Note that a regular covalent bond is denoted by a minus sign (−) and a hydrogen bond by three dots (···). The tolerances σ_c will be described in Section 4.

W3 Finally, the Balanced Loss function is minimized, and the R_c values of the optimal parameters are analyzed. When one category keeps dominating the loss function throughout the optimization, the only possible explanation is that the corresponding tolerances σ_c were set too small. There is no way to lower R_c because L_{BL} already ignores all of the other categories. The only option left is to accept that the model cannot reproduce items in category c with the desired accuracy, and to adjust one's expectations by increasing the corresponding tolerances σ_c. One can now repeat the parameter optimization and re-evaluate the result, possibly repeating the exercise a few times. Unlike tuning the weights in a conventional loss function, these repeated optimizations provide insight: they inform the human operator about the capabilities of the model and help manage expectations. It may also happen that some R_c end up well below 1, in which case we do not recommend decreasing the corresponding σ_c. Such a fortuitous outcome should not affect the desired accuracy.

Note that steps W2 and W3 in the above workflow involve (possibly subjective) human decisions and therefore cannot be replaced by an autonomous algorithm. This is an unavoidable aspect of multiobjective problems: one has to decide on a compromise that must be reached between different categories. The overall goal of Balanced Loss is to facilitate the finding of suitable compromises.

4. RESULTS AND DISCUSSION

4.1. Balanced Loss Optimization Procedure. This section illustrates how the optimization workflow of Balanced Loss leads to a competitive ReaxFF parametrization, using the alumina training set as a realistic example. In addition to the final force field, the intermediate steps provide insight into the capabilities of ReaxFF.

The data in the training set, described in Section 2.2, have been grouped into categories as described in Section 3 and as summarized in the leftmost column of Table 2. The optimizations are carried out in two stages, initial and final, which differ only in the tolerances σ_c.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

Table 2. Tolerances Used in the Balanced Loss Optimization

category	unit	initial σ_c	final σ_c
Al−O	Å	0.05	0.07
Al−Al	Å	0.05	0.10
O−H	Å	0.05	0.05
O···H	Å	0.05	0.12
Al−O−Al	deg	2.0	5.0
Al−O−H	deg	2.0	5.0
O−Al−O	deg	2.0	5.0
Al−O···H	deg	2.0	7.0
H−O···H	deg	2.0	7.0
H···O···H	deg	2.0	7.0
BSH	kcal mol^{-1}	1.25	4.0
GEH	kcal mol^{-1}	1.25	4.0
GSH	kcal mol^{-1}	1.25	4.0
SUR	kcal mol^{-1}	1.25	3.0
FOR	kcal mol^{-1}	1.25	3.0

Each entry is given an initial chemically relevant tolerance, σ_c, equal to the default sigma value from ParAMS, as shown in Table 2. With these tolerances, the ReaxFF parameters were optimized 40 times by using different random seeds to produce independent solutions. Figure 2a shows the evolution of the Balanced Loss during the 40 CMA runs. Figure 2c presents the $R_c\sigma_c$ values (category RMSEs with units) for the 40 optimized parameter vectors. The curves and data points are colored according to the loss value of the best parameter vector of each run. Of the 40 CMA runs in the initial stage, two are clearly worse than all others, presumably converging to unfavorable local minima. All 38 remaining runs produce comparable RMSE values but are not identical, which is the expected behavior. ReaxFF loss functions are known to exhibit many local minima and apparent noise due to the high sensitivity of the geometry optimizations in the training set to the ReaxFF parameters.[27] Adsorption energies in the categories **BSH** and **GSH** are the highest relative to their tolerance, σ_c. The performances in all other categories have less effect on the optimized parameters simply because these errors are closer to their tolerance. This means that, in this initial stage, the CMA runs train almost exclusively on the adsorption energies. It is therefore highly unlikely to find ReaxFF parameters that can further lower RMSE on the adsorption energies, let alone reach the tolerance of 1.25 kcal mol^{-1}. Also, for all other categories, the initial tolerances seem too optimistic, which will be addressed in the next stage.

For the second (and final) stage, the tolerances are revised, as shown in the last column of Table 2, to be more consistent with what ReaxFF could achieve in the initial stage. Without Balanced Loss, these tolerances can only be set by expert judgment, which is greatly facilitated here by the feedback from the initial stage. Figure 2b shows that the revised tolerances result in more consistent loss values across all 40 parametrizations. This is also reflected in the RMSEs in Figure 2d, most of which exhibit less scatter. In other words, for most categories, the 40 parametrizations in the final stage are comparable, with the category **GEH** being the most notable exception. Note that the absolute values of Balanced Loss cannot be compared between Figure 2a and b because the two optimization stages use different tolerances. One could slightly tweak the tolerances further to bring them closer to the errors on the training set, but this would amount to relatively small adjustments that we do not expect to lead to significant improvements.

To illustrate the importance of the LSE form of Balanced Loss, we have performed the 40 CMA optimizations in six different ways: with initial and final tolerances and using different functions f: $f(x) = \exp(x)$ (Log-Sum-Exp or LSE, the default, same results as above), $f(x) = x^2$ (root-sum-square or RSS), and $f(x) = x$ (identity-sum-identity or ISI). For comparison, the LSE form of Balanced Loss is computed for all 240 optimized parameter vectors, and their distribution is shown in Figure 3a. In the initial stage, function f has a significant influence. The choice of function f determines the compromise between the RMSEs of the individual categories: in the case of RSS and ISI, the parametrization no longer exclusively prioritizes the adsorption energies in the categories **BSH** and **GSH**, resulting in higher RMSEs for these categories, as illustrated in Figure 3b. The results in this figure do not reveal whether the poor performance of ReaxFF for the categories **BSH** and **GSH** can be remedied by giving these categories a higher weight in the loss function or whether they are high due to intrinsic limitations of the model. The LSE form automatically resolves this ambiguity. Since this form approximates the maximum over all $R_c\sigma_c$ this loss function is dominated by the categories **BSH** and **GSH**, resulting in the lowest possible RMSE for these categories. It is simply impossible to give these categories a higher weight, leaving no other option but to assign more humble tolerances.

Figure 3a reveals two additional insights. First, the results become less sensitive to the choice of function f in the final stage. The RMSEs for all categories are close to the final tolerances, meaning that the argument of the function f (for the optimized parameters) is close to one, reducing the importance of the nonlinearity of f. Second, with $f(x) = x^2$ (RSS), the loss function is mathematically equivalent to eq 5, a standard loss function used for ReaxFF. This implies that the optimized parameters in this work can also be found with a more conventional loss function when the weights and sigmas are set consistently with the tolerances in Balanced Loss. Hence, the added value of Balanced Loss is essentially the feedback provided from the initial stage, which facilitates the configuration of the tolerances.

Figure 4 shows the distribution of the 40 optimized parameter vectors after transforming them to their dimensionless form. Although the 40 CMA runs converge to approximately the same Balanced Loss value, the corresponding parameters are not necessarily similar. Some parameters, such as `Al.O:D_e^sigma` or `Al.O:p_ovun1`, have a delineated range of optimal values. However, most parameters can be found across the entire interval of the allowed values. This does not mean that all of these parameters are completely random: they could be correlated, which is not apparent in the individual histograms. In any case, the optimal parameters are degenerate to some degree, which has also been observed in previous ReaxFF parametrizations on other chemical systems.[21,27,31,44]

Since ReaxFF is at least partially inspired by physical principles, one might hope that all parameters always converge to the same values. However, Sethna *et al.* have extensively shown in their work on "sloppy models" that broad parameter distributions are virtually always found for models with more than a few parameters, across different scientific disciplines.[127−130] This is a universal pattern, regardless of the degree of physical interpretation the model parameters may

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

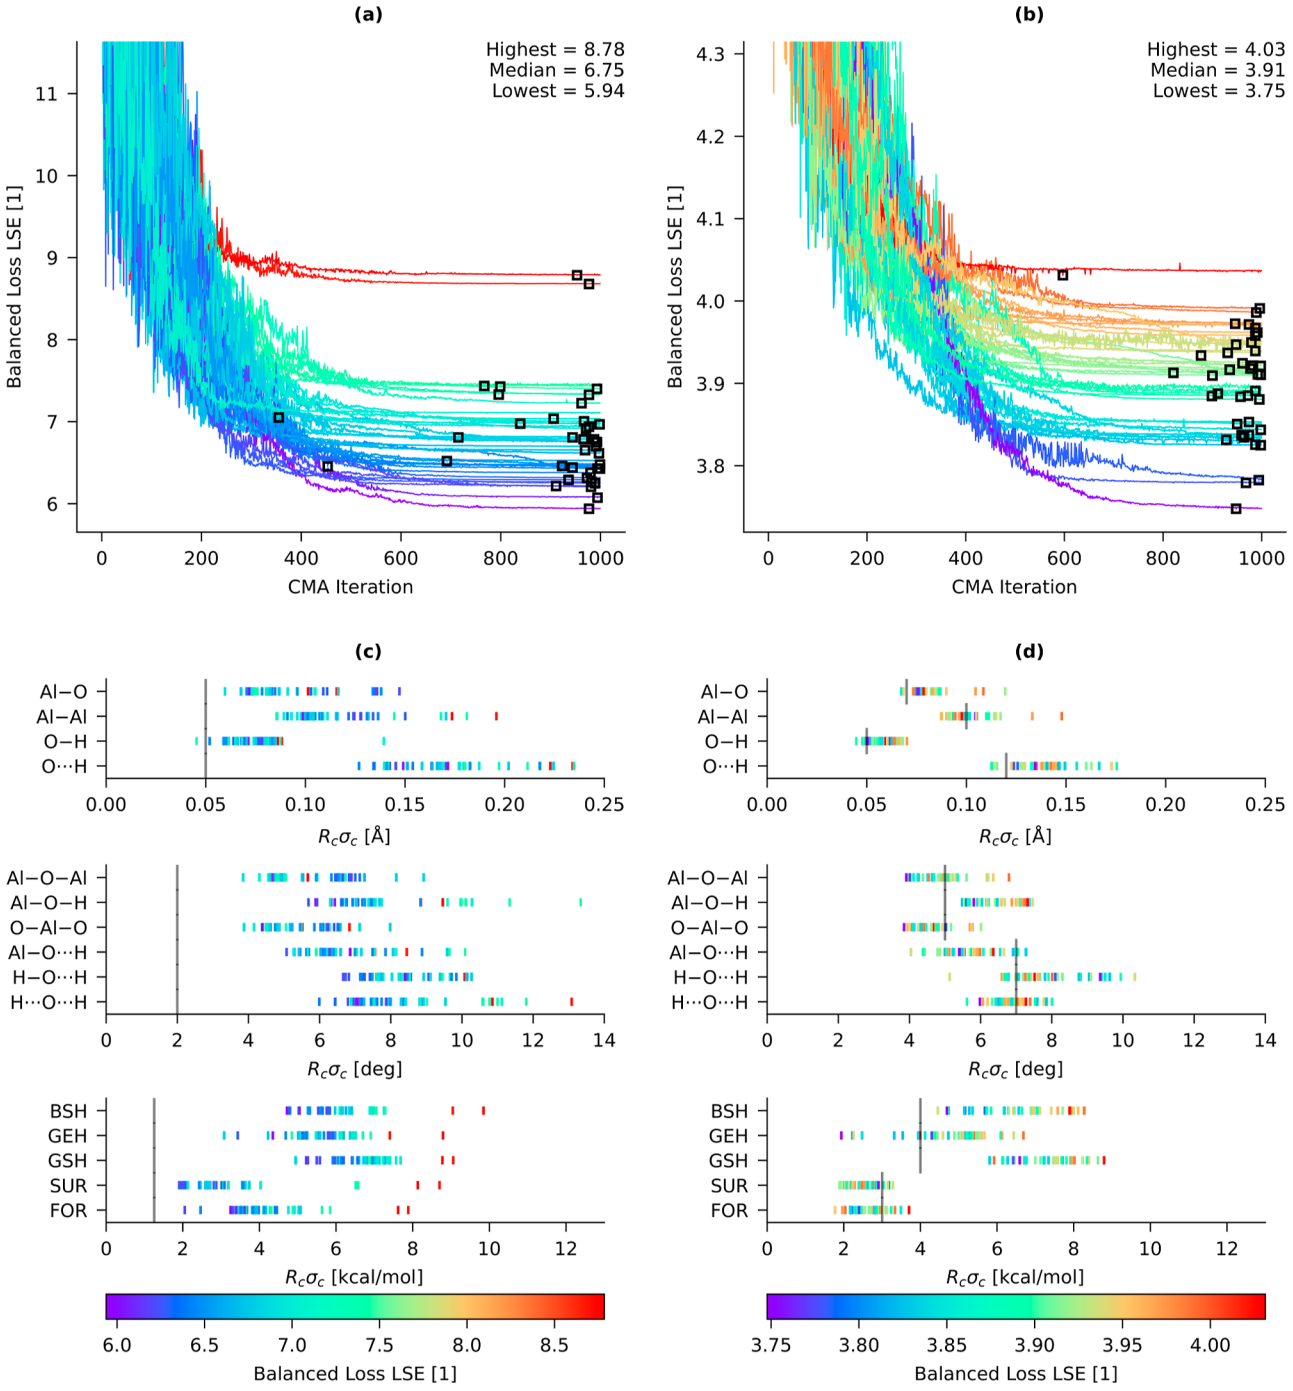

Figure 2. Balanced Loss as a function of CMA iteration (lowest value within the population) of the 40 parametrizations in the initial (a) and final (b) stages. Lowest value along each trajectory is indicated by a black square. Dimensioned category RMSEs, $R_c\sigma_c$, of the best parameter vector of each of the 40 parametrizations in the initial (c) and final (d) stage. All data are color coded by the loss value of the best point along the trajectory. Color bar of panel (c)/(d) is also applicable to panel (a)/(b). Gray vertical lines in panel (c)/(d) denote the tolerances for the corresponding category.

have. It is observed that some degrees of freedom in the parameter space of complex models are systematically ill-defined, not due to a lack of data but because nearly the same model predictions (for all possible inputs) are found for different parameter vectors. As a consequence, predictions based on unseen data are robust despite uncertainties in the parameters. For the electronegativity equalization method (EEM), which is included in ReaxFF, this parameter degeneracy has been investigated in more detail.[131] ReaxFF also has the characteristics of "a sloppy model", as illustrated by Figure S8 in the Supporting Information. All of the optimized parameter vectors of the 40 CMA runs (LSE, final stage) perform reasonably well for the validation set, despite the fact that they represent different local minima in the parameter space. This in itself is not a limitation when the model is used for simulations, but it obviously makes any direct interpretation of the parameters impossible.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

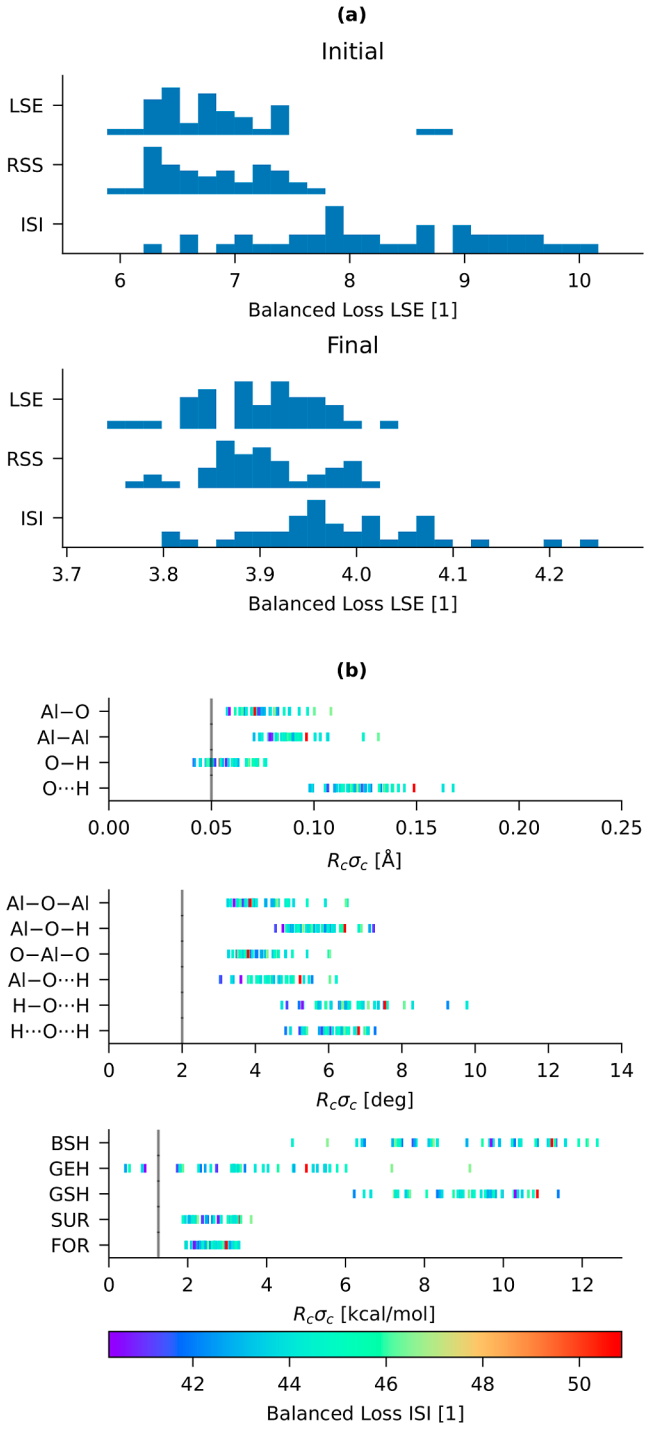

Figure 3. (a) Distribution of Balanced Loss values, computed with $f(x) = \exp(x)$, for parameter vectors optimized with different functions f in eq 7: $f(x) = \exp(x)$ (LSE), $f(x) = x^2$ (RSS), and $f(x) = x$ (ISI). Histograms are computed for both the initial and final tolerances defined in Table 2. (b) Dimensioned category RMSEs, $R_c\sigma_c$, of the best parameter vector of each of the 40 parametrizations in the initial stage using $f(x) = x$ (ISI).

It is also noticeable that several parameters have a high probability of converging close to the bounds. One might deduce that the parameter bounds are too narrow and the optimizer is trying to move the parameters to an optimum beyond the bounds, but this is not the only possible cause. Note that the components of the best parameter vector over all

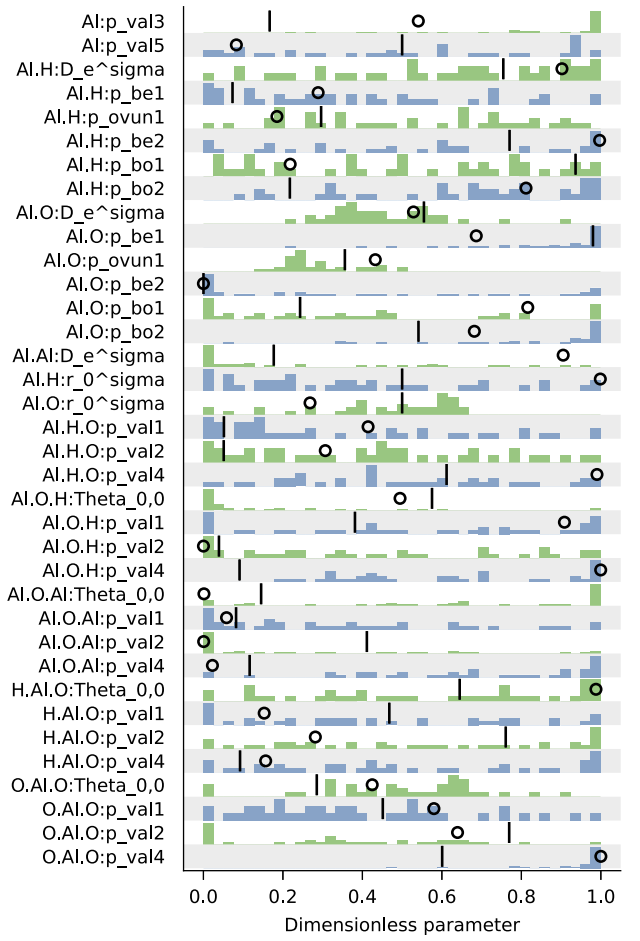

Figure 4. Histograms of the 36 components of the 40 optimized dimensionless parameter vectors in the final stage. Parameter components are made dimensionless by a linear transformation, such that zero corresponds to the lower bound and one corresponds to the upper bound. Bounds are listed in Table S3 in the Supporting Information. Initial values are marked with vertical black bars. Values corresponding to the lowest loss (over all 40 runs) are marked with circles.

40 runs, the circles in Figure 4, are not necessarily close to the bounds, even if the remaining near-optimal values cluster near the edges. Examples of this pattern are `Al.O:p_be1` and `Al.O.H:Theta_0,0`. A deeper investigation, beyond the scope of the current work, is needed to understand why a disproportionate number of near-optimal solutions are found near the bounds. For example, this could also be related to an optimizer inefficiency when the parameters approach their bounds, and addressing this problem may make the optimization more efficient. It should also be noted that some initial parameter values start close to the interval bounds regardless of the boundary extension, as explained in Section 2.3. Since some of the initial parameters have values close to zero, the effect of the boundary extension is negligible. The most prominent examples are `Al.O:p_be2`, `Al.H.O:p_val1`, `Al.H.O:p_val2`, and `Al.O.H:p_val2`.

For the remainder of this work, the best parameter vector from the final stage is used for all calculations, *i.e.*, corresponding to the lowest square in Figure 2b and the circles in Figure 4. The selected parameters are given in the last column of Table S3 in the Supporting Information.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

4.2. Force Field Validation. Figure 5 offers a first visual impression of how the new force field improves the prediction

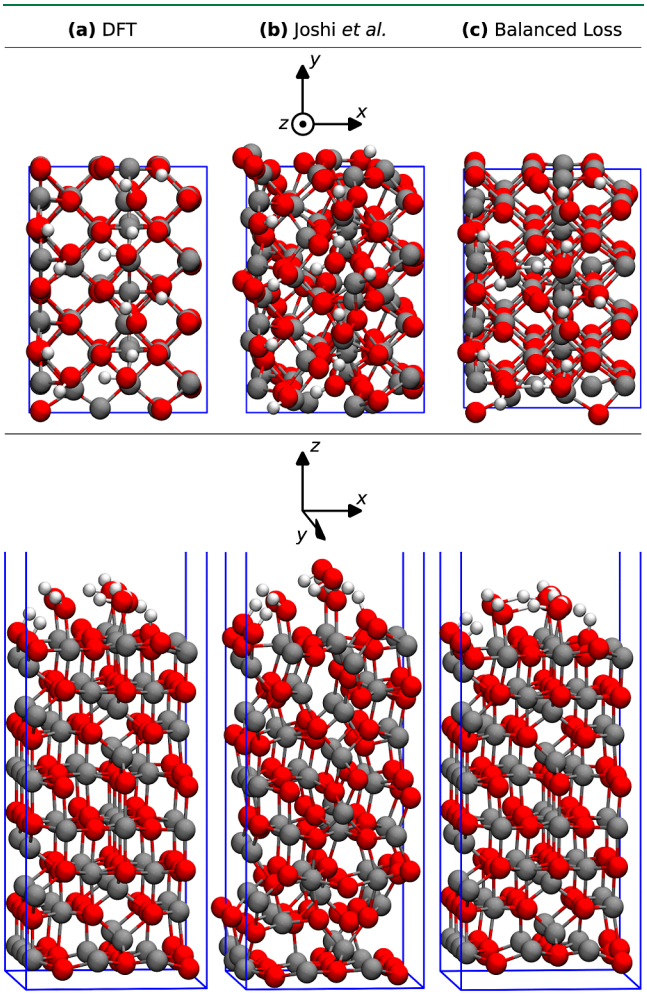

Figure 5. Optimized structures of the γ-Al_2O_3(001) surface from the validation set at a OH coverage of 13.0 nm^{-2}, computed with different models: (a) DFT reference, (b) ReaxFF parameters by Joshi *et al.*, and (c) ReaxFF parameters obtained with Balanced Loss. Al = gray, O = red, and H = white.

of dissociative water adsorption on alumina. It shows a high-hydration structure of the γ-Al_2O_3(001) slab model, labeled `gamma_surf-001_03w` in the validation set, geometry-optimized with the reference method (DFT calculation with VASP), the initial force field by Joshi *et al.*,[90] and the new force field in this work. The Joshi *et al.* force field predicts a severe deformation of the alumina structure and already desorbs water molecules in this static calculation, while all water should be adsorbed according to the reference calculation. In contrast, the new force field in this work predicts a geometry that is visually deviating only slightly from the DFT reference, and the water molecules remain adsorbed on the surface.

The performance of the new force field can be evaluated more in detail by analyzing the distributions of categorized data in the training and validation sets and their deviations from the DFT reference. Table 3 shows the RMSEs between force field and reference data, for each category, for the training and validation set, and for the initial force field by Joshi *et al.*[88,90] and the one optimized in this work. To make the RMSEs directly comparable, only data related to γ-Al_2O_3 was taken from the training set since the validation set also contains only γ-Al_2O_3. In addition, only categories (rows) present in both datasets are considered. For reference, the standard deviations on the reference data per category are also reported. Figure 6 shows parity plots for all categories of internal coordinates in the validation set. In addition, water adsorption energies on γ-Al_2O_3 surfaces in the training and validation sets are shown in Figure 7. Not all data categories from the training set exist for the validation set because the validation set is focused on adsorption on γ-Al_2O_3 only. This is reflected in Table 3 and Figures 6 and 7 by only considering γ-Al_2O_3 structures. Table S6 and Figure S3 in the Supporting Information contain the results omitted here, *i.e.*, not involving γ-Al_2O_3 surfaces, for which a direct comparison to the validation set is not possible.

The RMSEs in Table 3 show that the new force field significantly reduces the errors on the bond lengths compared to the Joshi *et al.* force field.[88,90] A subset of the bonds is broken after geometry optimization with the force fields, which is not fully visible in Figure 6a,c because this would require an impractical scale for the vertical axes. In the validation set, 14.3% of the Al−O bonds and 1.1% of the O−H bonds are broken with the Joshi *et al.* force field. With the new force field proposed here, these percentages reduce to 0.0 and 0.3%, respectively. These percentages are consistent with the visualization in Figure 5b and confirm that the Joshi γ-Al_2O_3 surface force field cannot preserve the structural integrity of the (hydrated) γ-Al_2O_3 slabs. By consequence, this force field also performs poorly for other categories (angles and energies), for which a correct bonding topology is required.

The parity plots of the distances in Figure 6a,c, and d also reveal that even the new force field captures the variations in bond lengths only approximately. This is also confirmed by the fact that the RMSEs of the distances in Table 3 are of the same order as the standard deviation on the distances in the reference data. For the O−H and O···H distances, this was to be expected because the corresponding bond parameters in ReaxFF were not reoptimized for the sake of compatibility with the silicate parameters in the Joshi *et al.* force field. For the Al−O distances, the performance is slightly better, which is consistent with the fact that several Al−O parameters were reoptimized.

The new force field also improves upon the Joshi *et al.* force field in terms of valence angles, again with a somewhat better performance when no hydrogen atoms are involved. It is remarkable that the multimodal distributions of the Al−O−Al, O−Al−O, and Al−O···H angles are reproduced well by the new force field, despite only having a single energy term for these angles in the ReaxFF force field. Also, the improvements of the H−O···H and H···O···H angles, compared to the initial force field, are remarkable because no corresponding valence angle or hydrogen bonding terms were reparameterized.

In line with the previous categories, the RMSEs on the energies are significantly smaller with the new force field. Before discussing the adsorption energies, it should be noted that the category **SUR** was mainly introduced to improve the diversity of the training set. (It comprises reaction energies between bulk and slab models, normalized on the number of Al atoms.) A few data points in the same category can be derived from the validation structures and are included here for the sake of completeness. The error on these data points has also decreased compared to the initial force field, confirming

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

Table 3. Comparison of Root-Mean-Square-Errors (RMSEs) of the Initial Force Field by Joshi *et al.*[88,90] and the Force Field Optimized in This Work (BL)[a]

category	unit	training				validation			
		ref	Joshi	BL		ref	Joshi	BL	
		SD	RMSE	RMSE	#	SD	RMSE	RMSE	#
Al−O	Å	0.16	0.45	0.09	1405	0.10	0.40	0.06	2102
Al−Al	Å	0.30	0.40	0.11	6759	0.31	0.40	0.11	13,588
O−H	Å	0.03	0.31	0.07	352	0.02	0.18	0.06	704
O···H	Å	0.18	0.85	0.17	203	0.18	0.90	0.25	351
Al−O−Al	deg	21.2	17.0	4.4	1756	20.0	16.6	3.8	2744
Al−O−H	deg	12.2	20.4	6.9	598	9.3	18.7	8.3	1217
O−Al−O	deg	29.4	16.7	4.5	2943	30.9	17.4	3.9	4656
Al−O···H	deg	18.4	16.9	6.5	301	17.6	17.6	6.6	554
H−O···H	deg	17.5	27.4	12.7	142	17.3	29.4	14.7	254
H···O···H	deg	31.1	28.9	4.4	65	38.0	33.7	7.7	88
GSH	kcal mol^{-1}	15.0	55.7	6.6	37	16.4	54.4	8.3	101
SUR	kcal mol^{-1}	6.1	11.3	2.9	8	1.4	4.7	1.4	6

[a]The RMSEs are computed for training and validation sets and are grouped per data category. For reference, the standard deviation (SD) on the reference data is included.

the ability of Balanced Loss to account for under-represented categories in the training set.

The new force field reduces the error on the adsorption energies (category **GSH**) by more than a factor of 6 in the validation set. The improvements are also immediately clear in Figure 7, which shows the adsorption energy on γ-Al_2O_3 surfaces as a function of the OH coverage. In most cases, the new force field predicts the correct trends in the adsorption energy, with some exceptions at low OH coverage in Figure 7d,f. For the (001) and $(110)_b$ slabs in the validation set, it is unclear why these surfaces exhibit larger errors in adsorption energy at low coverage. The Joshi *et al.* force field incorrectly predicts water desorption for 12 out of 53 surface structures in the validation set, which hampers reliable energy predictions. With, the new force field, this problem is far less prevalent: only one water molecule (out of six) from only one surface structure spontaneously desorbs. The new force field can also predict the magnitude of the adsorption energy, with a RMSE of 8.3 kcal mol^{-1}, compared to a standard deviation of 16.4 kcal mol^{-1} of the adsorption energies in the validation set.

For all categories of data discussed above, the errors on static calculations are very similar for the training and validation sets, indicating that the significant improvements of the new ReaxFF force field generalize to structures not used for training. The errors on distances, valence angles, and energy differences are also comparable to those reported for previous ReaxFF models.[8,25,132−135]

Because the training set only includes equilibrium geometries, it should be tested to what extent the new parameters can also reproduce nonequilibrium energies. To this end, a constant-temperature DFT MD run was performed on structure `gamma_surf-110l_A1_06w` from the validation set using the same level of theory as the training data. An elevated temperature of 1000 K stimulates the desorption of water, which is observed during the first 200 fs. Section S2 of the Supporting Information presents a detailed comparison of the DFT and ReaxFF energies computed for snapshots from this trajectory. In summary, the instantaneous DFT adsorption energy computed with eq 3 as a function of time is reproduced qualitatively by the ReaxFF parameters obtained with Balanced Loss: the relative error of about 25% over the first 200 fs is comparable to the RMSE on the training set for the GSH category. Our new parameters also show a clear improvement compared to the ReaxFF energies obtained with the parameters of Joshi *et al.* The thermal energy fluctuations due to vibrations within the alumina slab are not well reproduced, which is expected, since no corresponding data was used for training.

Despite our improvements, it remains interesting to explore further refinements, *e.g.*, to further reduce errors in adsorption energies or to improve the vibrational states of alumina. One avenue is to activate more parameters during the training that are now fixed for the sake of backward compatibility. Giving up backward compatibility would only be useful when extending the chemical space of the training set to aluminosilicates and water, such that all parameters of the Joshi *et al.* force field can be reoptimized. However, this would be a daunting enterprise because the current training set size is already computationally demanding: a single CMA run in this work already took more than 24 h. In addition to the increased cost of the training set, more active parameters also imply more local minima and a slower convergence of CMA-ES, further exacerbating the computational cost. This avenue is therefore feasible only when one can drastically speed up the training of ReaxFF parameters. It is encouraging that efficiency gains were reported in recent publications, *e.g.*, by parallel optimization management[44] or by machine learning surrogates of the loss function.[28,136,137] One may also reduce the dimensionality of the parameter space through sensitivity analysis to speed up the CMA runs.[31]

As a final check of the new force field, an MD simulation is performed on a $(110)_l$ γ-Al_2O_3 slab with a cross-section of 6.0 nm,[2] a thickness of 1.8 nm, and surrounded by a vacuum layer of 5.7 nm wide. In the initial structure, the maximal number of water molecules is dissociatively adsorbed on both sides of the slab, such that the chemical formula is $Al_{504}O_{864}H_{216}$. The MD simulation employs a Nosé−Hoover thermostat[138] with a time step of 0.2 fs, a temperature of 500 K, a time-constant of 500 ps, and a fixed periodic box size. Figure 8 shows the initial and final states of the MD trajectory, as well as the evolution of the kinetic energy, total energy (kinetic + potential), and conserved quantity. The conserved quantity (green) exhibits a slow linear increase, which is acceptable for long ReaxFF MD

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

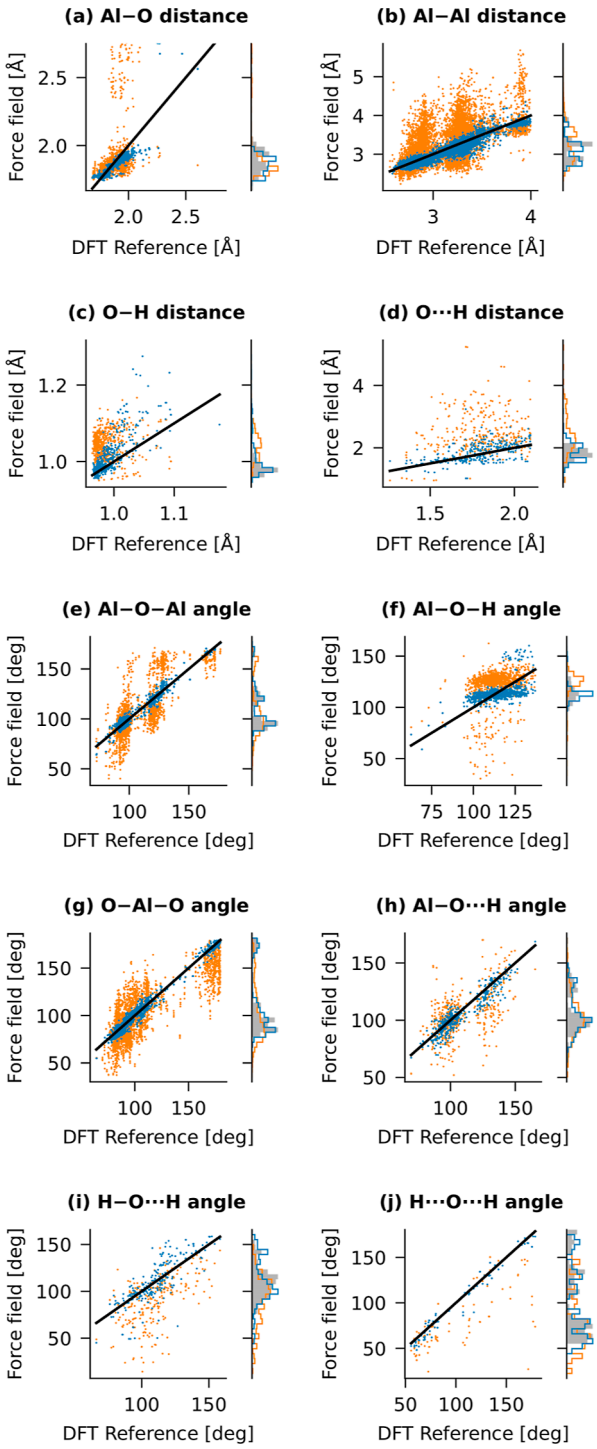

Figure 6. Parity plots of the reference internal coordinates from the validation set versus the force field predictions: initial parameters by Joshi *et al.*[88,90] (orange) and optimal parameters obtained with Balanced Loss (blue). Distances: (a) Al−O, (b) Al−Al, (c) O−H, and (d) O···H. Angles: (e) Al−O−Al, (f) Al−O−H, (g) O−Al−O, (h) Al−O···H, (i) Al−O···H, and (j) Al···O···H. The parity line is plotted as a black solid line. In panels (a,c), the vertical axis is manually limited to only show bonded distances. Initial parameters from Joshi result in many broken bonds, which are omitted for the sake of clarity.

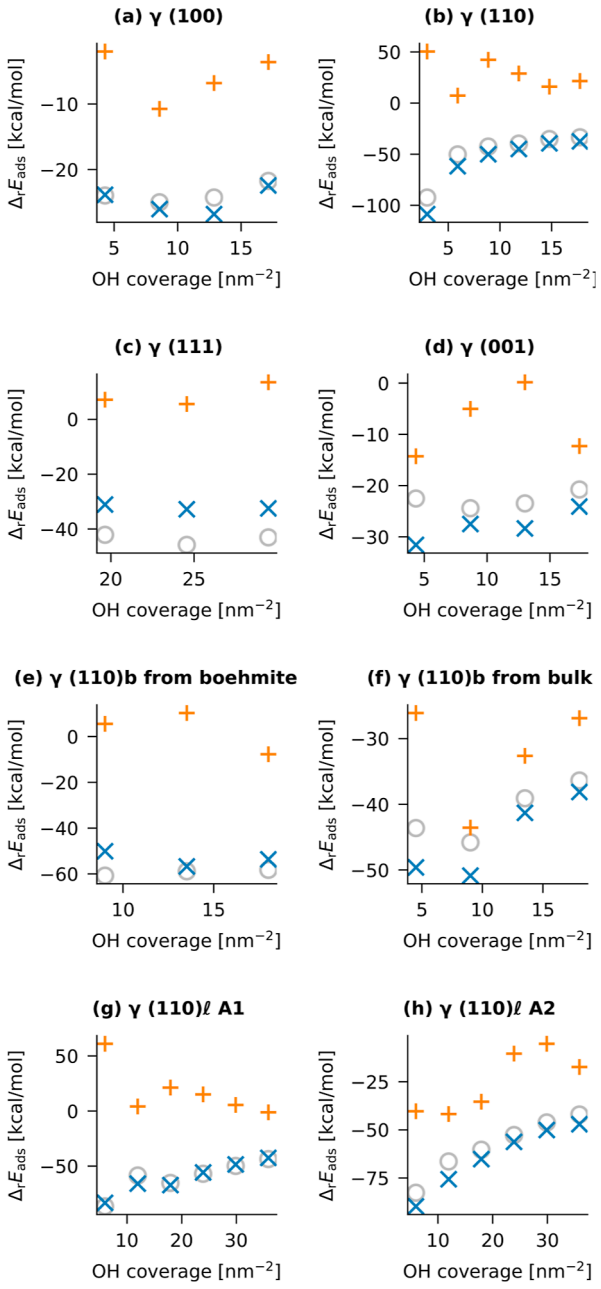

Figure 7. Water adsorption energy per water molecule on γ-Al_2O_3 surfaces as a function of the OH coverage, using the bare surface as the reference, *i.e.*, using $m_i = 0$ in eq 3. Results are computed with DFT (gray circle), Joshi FF (orange plus), and FF from this work (blue cross). Panels (a−c) are adsorption energies from the training set, whereas (d−h) are derived from the validation set. Nomenclature of the surfaces in the validation set is described in ref 52.

simulations at a constant temperature. ReaxFF forces are imperfect due to the numerical convergence of the variable charges and small discontinuities in the ReaxFF energy surface. Such small force errors are practically tiny random kicks on the nuclei, which slowly pump energy into the system, but this is easily compensated for by the thermostat and results in a slowly increasing conserved quantity. The 3D visualizations show that some of the water molecules desorb, as expected at a temperature of 500 K. This test shows that the new force field can also be used for MD simulations, even though it is only trained on optimized geometries. A complete study of water adsorption, with larger slabs, different alumina surfaces, and temperatures, goes beyond the scope of this work. We expect that the new force field will make such simulations possible at time and length scales that are infeasible for DFT methods.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

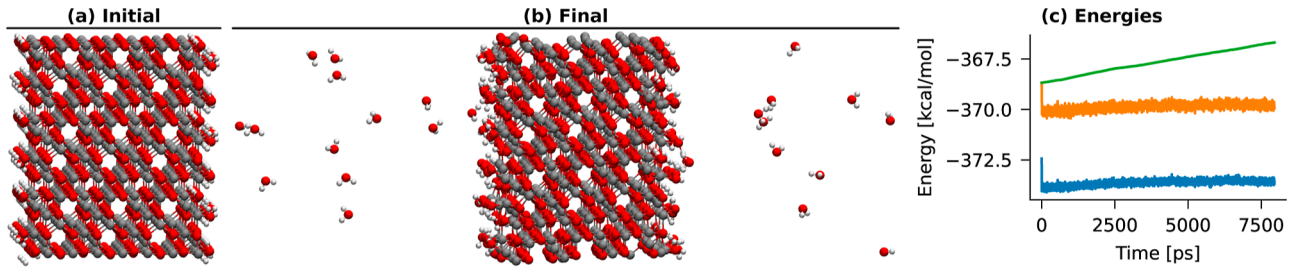

Figure 8. Initial (a) and final (b) states of a constant-temperature MD trajectory of a $(110)_l$ γ-Al_2O_3 with water initially adsorbed on the surface. Al = gray, O = red, and H = white. See text for details. (c) Kinetic energy (blue), total energy (kinetic + potential, orange), and the conserved quantity (green).

5. CONCLUSIONS AND OUTLOOK

This work addresses the difficulty of assigning fitting weights (or their inverses, often called sigmas) in a conventional ReaxFF loss function. Balanced Loss is proposed as a new cost function as well as a workflow to reformulate the weight assignment in terms of more manageable concepts. One starts by classifying the training data into meaningful categories with a corresponding tolerance, which is the RMSE (between ReaxFF predictions and reference data in that category) that one is willing to tolerate. When the error on one category exceeds the corresponding tolerance more than on other categories, the LSE form of Balanced Loss guarantees that this error will completely dominate the loss function, effectively forcing the optimizer to reduce this error first. If one or more of these dominating categories remain after the parameters converge, it is guaranteed that these errors cannot be reduced further at the expense of making larger errors in other categories. As a result, parameter optimization also assesses whether ReaxFF can meet the expectations defined by the tolerances. If necessary, the expectations can be adjusted, followed by a new parameter optimization. The methodology is applied to the realistic and challenging reparameterization of ReaxFF for water adsorption on alumina surfaces. This not only results in a competitive force field but also provides insight into the performance that can be expected from ReaxFF for each category of training data.

The new force field derived in this work is a refinement of the alumina parameters in the aluminosilicate force field of Joshi *et al.*[88,90] The training and validation data consisted of geometry and energy data from previous DFT studies of water adsorption on boehmite and γ-Al_2O_3 surfaces. While γ-Al_2O_3 is industrially the most relevant, boehmite structures were included in the training set to improve the data diversity. Parameter selection focused on maintaining backward compatibility with the ReaxFF model of Joshi *et al.* as much as possible while still activating sufficient parameters to reproduce the training data. CMA is used to minimize Balanced Loss as a function of the selected parameters. 40 independent CMA runs were performed to test the robustness of the optimized parameters. Of all these runs, the result with the lowest error on the training set is used for validation. Static calculations confirm that the optimized force field produces very similar errors on the γ-Al_2O_3 properties present in the training and validation sets. The force field can be used for MD simulations, and we expect it to be applicable to extensive simulations of water adsorption on alumina surfaces at time and length scales inaccessible to DFT methods.

This study also revealed several challenges and new avenues for future method development. Obviously, it should be validated whether the Balanced Loss workflow is equally helpful for other ReaxFF parametrizations. Even beyond the scope of ReaxFF, Balanced Loss may facilitate optimization problems involving multiple (possibly competing) categories of training data. In addition, this study confirmed the known pitfalls of ReaxFF parametrization and suggested new ones. Despite our careful and relatively small selection of parameters, the minimum of the loss function is still degenerate, suggesting that the number of active parameters can be reduced further with sensitivity analysis. Finally, it was observed that parameters often converge to near-optimal values close to the bounds. A better understanding of this phenomenon may help to speed up convergence to better parameters.

■ ASSOCIATED CONTENT

Data Availability Statement

A dataset as made available at 10.5281/zenodo.10491516 comprising: the training and validation data, the results of the 240 CMA optimizations, and scripts used to select parameters, convert datasets, and perform the training and validation.

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jctc.3c01009.

Training and validation sets, ReaxFF parameters selected for optimization, and additional performance metrics on the training set that have no counterpart in the validation set (PDF)

Optimized parameters used for the force field validation and example input files for the MD simulations (ZIP)

■ AUTHOR INFORMATION

Corresponding Author

Toon Verstraelen − *Center for Molecular Modeling (CMM), Ghent University, B-9052 Ghent, Belgium;* orcid.org/0000-0001-9288-5608; Email: toon.verstraelen@ugent.be

Authors

Loïc Dumortier − *IFP Energies nouvelles, 92852 Rueil-Malmaison, France; Center for Molecular Modeling (CMM), Ghent University, B-9052 Ghent, Belgium*

Céline Chizallet − *IFP Energies nouvelles, 69360 Solaize, France;* orcid.org/0000-0001-5140-8397

Benoit Creton − *IFP Energies nouvelles, 92852 Rueil-Malmaison, France;* orcid.org/0000-0002-3287-877X

Theodorus de Bruin − *IFP Energies nouvelles, 92852 Rueil-Malmaison, France;* orcid.org/0000-0001-6997-1786

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jctc.3c01009

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

The authors thank the company Software for Chemistry & Materials B.V. (SCM) for providing an AMS developer license at no cost. In particular, we are grateful to Matti Hellström and Tomáš Trnka from SCM for their technical support. The authors also thank members of the Center for Molecular Modeling at Ghent University, especially Tom Braeckevelt, Leonid Komissarov, and Michael Gustavo, as well as Pascal Raybaud from IFP Energies nouvelles for their insightful discussions. Computational resources and services used were provided by Ghent University (Stevin and Hortense Supercomputer Infrastructures), the VSC (Flemish Supercomputer Center), funded by the Research Foundation-Flanders (FWO), and IFP Energies nouvelles (ENER440). Funding was provided by IFPEN and UGent (BOF/24J/2023/121). T.V. acknowledges the Special Research Fund (BOF) of Ghent University for its continuous support.

■ REFERENCES

(1) Brunk, E.; Rothlisberger, U. Mixed Quantum Mechanical/Molecular Mechanical Molecular Dynamics Simulations of Biological Systems in Ground and Electronically Excited States. *Chem. Rev.* **2015**, *115*, 6217−6263.

(2) Döntgen, M.; Przybylski-Freund, M.-D.; Kröger, L. C.; Kopp, W. A.; Ismail, A. E.; Leonhard, K. Automated Discovery of Reaction Pathways, Rate Constants, and Transition States Using Reactive Molecular Dynamics Simulations. *J. Chem. Theory Comput.* **2015**, *11*, 2517−2524.

(3) Chenoweth, K.; van Duin, A. C. T.; Dasgupta, S.; Goddard III, W. A. Initiation Mechanisms and Kinetics of Pyrolysis and Combustion of JP-10 Hydrocarbon Jet Fuel. *J. Phys. Chem. A* **2009**, *113*, 1740−1746.

(4) Ding, J.; Zhang, L.; Zhang, Y.; Han, K.-L. A Reactive Molecular Dynamics Study of *n*-Heptane Pyrolysis at High Temperature. *J. Phys. Chem. A* **2013**, *117*, 3266−3278.

(5) Li, X.; Zheng, M.; Ren, C.; Guo, L. ReaxFF Molecular Dynamics Simulations of Thermal Reactivity of Various Fuels in Pyrolysis and Combustion. *Energy Fuels* **2021**, *35*, 11707−11739.

(6) Fogarty, J. C.; Aktulga, H. M.; Grama, A. Y.; van Duin, A. C. T.; Pandit, S. A. A reactive molecular dynamics simulation of the silica-water interface. *J. Chem. Phys.* **2010**, *132*, 174704.

(7) Mueller, J. E.; van Duin, A. C. T.; Goddard, W. A. Application of the ReaxFF Reactive Force Field to Reactive Dynamics of Hydrocarbon Chemisorption and Decomposition. *J. Phys. Chem. C* **2010**, *114*, 5675−5685.

(8) Müller, J.; Hartke, B. reaxFF Reactive Force Field for Disulfide Mechanochemistry, Fitted to Multireference ab Initio Data. *J. Chem. Theory Comput.* **2016**, *12*, 3913−3925.

(9) Rimsza, J. M.; Jones, R. E.; Criscenti, L. J. Crack propagation in silica from reactive classical molecular dynamics simulations. *J. Am. Ceram. Soc.* **2018**, *101*, 1488−1499.

(10) Mao, Q.; van Duin, A. C.; Luo, K. Formation of incipient soot particles from polycyclic aromatic hydrocarbons: A ReaxFF molecular dynamics study. *Carbon* **2017**, *121*, 380−388.

(11) Lei, T.; Guo, W.; Liu, Q.; Jiao, H.; Cao, D.-B.; Teng, B.; Li, Y.-W.; Liu, X.; Wen, X.-D. Mechanism of Graphene Formation via Detonation Synthesis: A DFTB Nanoreactor Approach. *J. Chem. Theory Comput.* **2019**, *15*, 3654−3665.

(12) van Duin, A. C. T.; Dasgupta, S.; Lorant, F.; Goddard, W. A. ReaxFF: A Reactive Force Field for Hydrocarbons. *J. Phys. Chem. A* **2001**, *105*, 9396−9409.

(13) van Duin, A. C. T.; Strachan, A.; Stewman, S.; Zhang, Q.; Xu, X.; Goddard, W. A. $ReaxFF_{SiO}$ Reactive Force Field for Silicon and Silicon Oxide Systems. *J. Phys. Chem. A* **2003**, *107*, 3803−3811.

(14) Senftle, T. P.; Hong, S.; Islam, M. M.; Kylasa, S. B.; Zheng, Y.; Shin, Y. K.; Junkermeier, C.; Engel-Herbert, R.; Janik, M. J.; Aktulga, H. M.; et al. The ReaxFF reactive force-field: development, applications and future directions. *npj Comput. Mater.* **2016**, *2*, 15011.

(15) Tersoff, J. Empirical Interatomic Potential for Carbon, with Applications to Amorphous Carbon. *Phys. Rev. Lett.* **1988**, *61*, 2879−2882.

(16) O'Connor, T. C.; Andzelm, J.; Robbins, M. O. AIREBO-M: A reactive model for hydrocarbons at extreme pressures. *J. Chem. Phys.* **2015**, *142*, 024903.

(17) Zhang, D.; Fonseca, A. F.; Liang, T.; Phillpot, S. R.; Sinnott, S. B. Dynamics of graphene/Al interfaces using COMB3 potentials. *Phys. Rev. Mater.* **2019**, *3*, 114002.

(18) Rowe, P.; Deringer, V. L.; Gasparotto, P.; Csányi, G.; Michaelides, A. An accurate and transferable machine learning potential for carbon. *J. Chem. Phys.* **2020**, *153*, 034702.

(19) Xue, L.-Y.; Guo, F.; Wen, Y.-S.; Feng, S.-Q.; Huang, X.-N.; Guo, L.; Li, H.-S.; Cui, S.-X.; Zhang, G.-Q.; Wang, Q.-L. ReaxFF-MPNN machine learning potential: a combination of reactive force field and message passing neural networks. *Phys. Chem. Chem. Phys.* **2021**, *23*, 19457−19464.

(20) Schreiner, M.; Bhowmik, A.; Vegge, T.; Busk, J.; Winther, O. Transition1x - a dataset for building generalizable reactive machine learning potentials. *Sci. Data* **2022**, *9*, 779.

(21) Pahari, P.; Chaturvedi, S. Determination of best-fit potential parameters for a reactive force field using a genetic algorithm. *J. Mol. Model.* **2012**, *18*, 1049−1061.

(22) Deetz, J. D.; Faller, R. Parallel Optimization of a Reactive Force Field for Polycondensation of Alkoxysilanes. *J. Phys. Chem. B* **2014**, *118*, 10966−10978.

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

(24) Dittner, M.; Müller, J.; Aktulga, H. M.; Hartke, B. Efficient global optimization of reactive force-field parameters. *J. Comput. Chem.* **2015**, *36*, 1550−1561.

(25) Trnka, T.; Tvaroška, I.; Koča, J. Automated Training of ReaxFF Reactive Force Fields for Energetics of Enzymatic Reactions. *J. Chem. Theory Comput.* **2018**, *14*, 291−302.

(26) Furman, D.; Carmeli, B.; Zeiri, Y.; Kosloff, R. Enhanced Particle Swarm Optimization Algorithm: Efficient Training of ReaxFF Reactive Force Fields. *J. Chem. Theory Comput.* **2018**, *14*, 3100−3112.

(27) Shchygol, G.; Yakovlev, A.; Trnka, T.; van Duin, A. C. T.; Verstraelen, T. ReaxFF Parameter Optimization with Monte-Carlo and Evolutionary Algorithms: Guidelines and Insights. *J. Chem. Theory Comput.* **2019**, *15*, 6799−6812.

(28) Sengul, M. Y.; Song, Y.; Nayir, N.; Gao, Y.; Hung, Y.; Dasgupta, T.; van Duin, A. C. T. INDEEDopt: a deep learning-based ReaxFF parameterization framework. *npj Comput. Mater.* **2021**, *7*, 68.

(29) Kaymak, M. C.; Rahnamoun, A.; O'Hearn, K. A.; van Duin, A. C. T.; Merz, K. M.; Aktulga, H. M. JAX-ReaxFF: A Gradient-Based Framework for Fast Optimization of Reactive Force Fields. *J. Chem. Theory Comput.* **2022**, *18*, 5181−5194.

(30) Komissarov, L.; Krep, L.; Schmalz, F.; Kopp, W. A.; Leonhard, K.; Verstraelen, T. A Reactive Molecular Dynamics Study of Chlorinated Organic Compounds. Part I: Force Field Development. *ChemPhysChem* **2023**, *24*, No. e202200786.

(31) Freitas Gustavo, M.; Hellström, M.; Verstraelen, T. Sensitivity Analysis for ReaxFF Reparametrization Using the Hilbert−Schmidt Independence Criterion. *J. Chem. Theory Comput.* **2023**, *19*, 2557−2573.

(32) Wang, H.; Zhang, L.; Han, J.; E, W. DeePMD-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Comput. Phys. Commun.* **2018**, *228*, 178−184.

(33) Unke, O. T.; Meuwly, M. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. *J. Chem. Theory Comput.* **2019**, *15*, 3678−3693.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

(34) Cools-Ceuppens, M.; Dambre, J.; Verstraelen, T. Modeling Electronic Response Properties with an Explicit-Electron Machine Learning Potential. *J. Chem. Theory Comput.* **2022**, *18*, 1672−1691.
(35) Fedkin, M. V.; Shin, Y. K.; Dasgupta, N.; Yeon, J.; Zhang, W.; van Duin, D.; van Duin, A. C. T.; Mori, K.; Fujiwara, A.; Machida, M.; et al. Development of the ReaxFF Methodology for Electrolyte−Water Systems. *J. Phys. Chem. A* **2019**, *123*, 2125−2141.
(36) Brown, I.; Smith, R.; Kenny, S. D. A ReaxFF potential for Al−ZnO systems. *Modell. Simul. Mater. Sci. Eng.* **2022**, *30*, 035001.
(37) Verstraelen, T.; Sukhomlinov, S. V.; Van Speybroeck, V.; Waroquier, M.; Smirnov, K. S. Computation of Charge Distribution and Electrostatic Potential in Silicates with the Use of Chemical Potential Equalization Models. *J. Phys. Chem. C* **2012**, *116*, 490−504.
(38) Bureekaew, S.; Amirjalayer, S.; Tafipolsky, M.; Spickermann, C.; Roy, T. K.; Schmid, R. MOF-FF − A flexible first-principles derived force field for metal-organic frameworks. *Phys. Status Solidi B* **2013**, *250*, 1128−1141.
(39) Grimme, S.; Bannwarth, C.; Shushkov, P. A Robust and Accurate Tight-Binding Quantum Chemical Method for Structures, Vibrational Frequencies, and Noncovalent Interactions of Large Molecular Systems Parametrized for All spd-Block Elements (Z = 1−86). *J. Chem. Theory Comput.* **2017**, *13*, 1989−2009.
(40) Komissarov, L.; Verstraelen, T. Improving the Silicon Interactions of GFN-xTB. *J. Chem. Inf. Model.* **2021**, *61*, 5931−5937.
(41) Li, J.; Song, X.; Li, P.; Herzfeld, J. A Carbon Is a Carbon Is a Carbon: Orbital-Free Simulations of Hydrocarbon Chemistry without Resort to Atom Types. *J. Phys. Chem. A* **2022**, *126*, 8468−8475.
(42) Włodarczyk, A.; Uchroński, M.; Podsiadły-Paszkowska, A.; Irek, J.; Szyja, B. M. Mixing ReaxFF parameters for transition metal oxides using force-matching method. *J. Mol. Model.* **2022**, *28*, 8.
(43) Komissarov, L.; Rüger, R.; Hellström, M.; Verstraelen, T. ParAMS: Parameter Optimization for Atomistic and Molecular Simulations. *J. Chem. Inf. Model.* **2021**, *61*, 3737−3743.
(44) Freitas Gustavo, M.; Verstraelen, T. GloMPO (Globally Managed Parallel Optimization): a tool for expensive, black-box optimizations, application to ReaxFF reparameterizations. *J. Cheminf.* **2022**, *14*, 7.
(45) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **2017**, *8*, 3192−3203.
(46) Chen, C.; Ong, S. P. A universal graph deep learning interatomic potential for the periodic table. *Nat. Comput. Sci.* **2022**, *2*, 718−728.
(47) Takamoto, S.; Shinagawa, C.; Motoki, D.; Nakago, K.; Li, W.; Kurata, I.; Watanabe, T.; Yayama, Y.; Iriguchi, H.; Asano, Y.; et al. Towards universal neural network potential for material discovery applicable to arbitrary combination of 45 elements. *Nat. Commun.* **2022**, *13*, 2991.
(48) Hart, L. D.; Lense, E. *Alumina Chemicals: Science and Technology Handbook*; John Wiley & Sons, 1990.
(49) Lefèvre, G.; Duc, M.; Lepeut, P.; Caplain, R.; Fédoroff, M. Hydration of γ-Alumina in Water and Its Effects on Surface Reactivity. *Langmuir* **2002**, *18*, 7530−7537.
(50) Arrouvel, C.; Digne, M.; Breysse, M.; Toulhoat, H.; Raybaud, P. Effects of morphology on surface hydroxyl concentration: a DFT comparison of anatase−TiO_2 and γ-alumina catalytic supports. *J. Catal.* **2004**, *222*, 152−166.
(51) Lagauche, M.; Larmier, K.; Jolimaitre, E.; Barthelet, K.; Chizallet, C.; Favergeon, L.; Pijolat, M. Thermodynamic Characterization of the Hydroxyl Group on the γ-Alumina Surface by the Energy Distribution Function. *J. Phys. Chem. C* **2017**, *121*, 16770−16782.
(52) Pigeon, T.; Chizallet, C.; Raybaud, P. Revisiting γ-alumina surface models through the topotactic transformation of boehmite surfaces. *J. Catal.* **2022**, *405*, 140−151.
(53) Levin, I.; Brandon, D. Metastable Alumina Polymorphs: Crystal Structures and Transition Sequences. *J. Am. Ceram. Soc.* **1998**, *81*, 1995−2012.
(54) Trueba, M.; Trasatti, S. P. γ-Alumina as a Support for Catalysts: A Review of Fundamental Aspects. *Eur. J. Inorg. Chem.* **2005**, *2005*, 3393−3403.
(55) Euzen, P.; Raybaud, P.; Krokidis, X.; Toulhoat, H.; Le Loarer, J.-L.; Jolivet, J.-P.; Froidefond, C. In *Handbook of Porous Solids*; Schüth, F., Sing, K. S. W., Weitkamp, J., Eds.; John Wiley & Sons, Ltd, 2002; pp 1591−1677.
(56) Knözinger, H. Dehydration of Alcohols on Aluminum Oxide. *Angew. Chem., Int. Ed. Engl.* **1968**, *7*, 791−805.
(57) Knözinger, H.; Bühl, H.; Kochloefl, K. The dehydration of alcohols on alumina: XIV. Reactivity and mechanism. *J. Catal.* **1972**, *24*, 57−68.
(58) Phung, T. K.; Lagazzo, A.; Rivero Crespo, M. A.; Sánchez Escribano, V.; Busca, G. A study of commercial transition aluminas and of their catalytic activity in the dehydration of ethanol. *J. Catal.* **2014**, *311*, 102−113.
(59) Kohl, A. L.; Nielsen, R. B. In *Gas Purification*, 5th ed.; Kohl, A. L., Nielsen, R. B., Eds.; Gulf Professional Publishing: Houston, 1997; pp 670−730.
(60) Larmier, K.; Chizallet, C.; Cadran, N.; Maury, S.; Abboud, J.; Lamic-Humblot, A.-F.; Marceau, E.; Lauron-Pernot, H. Mechanistic Investigation of Isopropanol Conversion on Alumina Catalysts: Location of Active Sites for Alkene/Ether Production. *ACS Catal.* **2015**, *5*, 4423−4437.
(61) Jain, J. R.; Pillai, C. N. Catalytic dehydration of alcohols over alumina: Mechanism of ether formation. *J. Catal.* **1967**, *9*, 322−330.
(62) Raybaud, P.; Toulhoat, H. *Catalysis by Transition Metal Sulphides: From Molecular Theory to Industrial Application*; Editions TECHNIP, 2013.
(63) Copéret, C.; Comas-Vives, A.; Conley, M. P.; Estes, D. P.; Fedorov, A.; Mougel, V.; Nagae, H.; Núñez-Zarur, F.; Zhizhko, P. A. Surface Organometallic and Coordination Chemistry toward Single-Site Heterogeneous Catalysts: Strategies, Methods, Structures, and Activities. *Chem. Rev.* **2016**, *116*, 323−421.
(64) Boudart, M. In *Advanced Catalysis*; Eley, D. D., Pines, H., Weisz, P. B., Eds.; Academic Press, 1969; pp 153−166.
(65) Prins, R. On the structure of γ-Al_2O_3. *J. Catal.* **2020**, *392*, 336−346.
(66) Valero, M. C.; Raybaud, P. Computational chemistry approaches for the preparation of supported catalysts: Progress and challenges. *J. Catal.* **2020**, *391*, 539−547.
(67) Krokidis, X.; Raybaud, P.; Gobichon, A.-E.; Rebours, B.; Euzen, P.; Toulhoat, H. Theoretical Study of the Dehydration Process of Boehmite to γ-Alumina. *J. Phys. Chem. B* **2001**, *105*, 5121−5130.
(68) Wischert, R.; Laurent, P.; Copéret, C.; Delbecq, F.; Sautet, P. γ-Alumina: the essential and unexpected role of water for the structure, stability, and reactivity of "defect" sites. *J. Am. Chem. Soc.* **2012**, *134*, 14430−14449.
(69) Digne, M.; Sautet, P.; Raybaud, P.; Euzen, P.; Toulhoat, H. Hydroxyl Groups on γ-Alumina Surfaces: A DFT Study. *J. Catal.* **2002**, *211*, 1−5.
(70) Digne, M.; Sautet, P.; Raybaud, P.; Euzen, P.; Toulhoat, H. Use of DFT to achieve a rational understanding of acid−basic properties of γ-alumina surfaces. *J. Catal.* **2004**, *226*, 54−68.
(71) Paglia, G.; Buckley, C. E.; Rohl, A. L.; Hart, R. D.; Winter, K.; Studer, A. J.; Hunter, B. A.; Hanna, J. V. Boehmite Derived γ-Alumina System. 1. Structural Evolution with Temperature, with the Identification and Structural Determination of a New Transition Phase, γ'-Alumina. *Chem. Mater.* **2004**, *16*, 220−236.
(72) Réocreux, R.; Girel, É; Clabaut, P.; Tuel, A.; Besson, M.; Chaumonnot, A.; Cabiac, A.; Sautet, P.; Michel, C. Reactivity of shape-controlled crystals and metadynamics simulations locate the weak spots of alumina in water. *Nat. Commun.* **2019**, *10*, 3139.
(73) Mardilovich, P. P.; Govyadinov, A. N.; Mukhurov, N. I.; Rzhevskii, A. M.; Paterson, R. New and modified anodic alumina membranes Part I. Thermotreatment of anodic alumina membranes. *J. Membr. Sci.* **1995**, *98*, 131−142.
(74) Brown, G. E.; Henrich, V. E.; Casey, W. H.; Clark, D. L.; Eggleston, C.; Felmy, A.; Goodman, D. W.; Grätzel, M.; Maciel, G.;

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

McCarthy, M. I.; et al. Metal Oxide Surfaces and Their Interactions with Aqueous Solutions and Microbial Organisms. *Chem. Rev.* **1999**, *99*, 77−174.
(75) Batista, A. T. F.; Wisser, D.; Pigeon, T.; Gajan, D.; Diehl, F.; Rivallan, M.; Catita, L.; Gay, A.-S.; Lesage, A.; Chizallet, C.; et al. Beyond γ-Al_2O_3 crystallite surfaces: The hidden features of edges revealed by solid-state 1H NMR and DFT calculations. *J. Catal.* **2019**, *378*, 140−143.
(76) Batista, A.; Pigeon, T.; Meyet, J.; Wisser, D.; Rivallan, M.; Gajan, D.; Catita, L.; Diehl, F.; Gay, A.; Chizallet, C.; et al. Structure, Location, and Spatial Proximities of Hydroxyls on γ-Alumina Crystallites by High-Resolution Solid-State NMR and DFT Modeling: Why Edges Hold the Key. *ACS Catal.* **2023**, *13*, 6536−6548.
(77) Pitman, M. C.; van Duin, A. C. T. Dynamics of Confined Reactive Water in Smectite Clay−Zeolite Composites. *J. Am. Chem. Soc.* **2012**, *134*, 3042−3053.
(78) Rimsza, J. M.; Yeon, J.; van Duin, A. C. T.; Du, J. Water Interactions with Nanoporous Silica: Comparison of ReaxFF and ab Initio based Molecular Dynamics Simulations. *J. Phys. Chem. C* **2016**, *120*, 24803−24816.
(79) Porter, A. J.; O'Malley, A. J. A Classical Molecular Dynamics Study on the Effect of Si/Al Ratio and Silanol Nest Defects on Water Diffusion in Zeolite HY. *J. Phys. Chem. C* **2021**, *125*, 11567−11579.
(80) van Gunsteren, W. F.; Berendsen, H. J. C. Computer Simulation of Molecular Dynamics: Methodology, Applications, and Perspectives in Chemistry. *Angew. Chem., Int. Ed. Engl.* **1990**, *29*, 992−1023.
(81) Nakata, A.; Baker, J. S.; Mujahed, S. Y.; Poulton, J. T. L.; Arapan, S.; Lin, J.; Raza, Z.; Yadav, S.; Truflandier, L.; Miyazaki, T.; et al. Large scale and linear scaling DFT with the CONQUEST code. *J. Chem. Phys.* **2020**, *152*, 164112.
(82) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khaliullin, R. Z.; Schütt, O.; Schiffmann, F.; et al. CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **2020**, *152*, 194103.
(83) Raybaud, P.; Digne, M.; Iftimie, R.; Wellens, W.; Euzen, P.; Toulhoat, H. Morphology and Surface Properties of Boehmite (γ-AlOOH): A Density Functional Theory Study. *J. Catal.* **2001**, *201*, 236−246.
(84) Motta, A.; Gaigeot, M.-P.; Costa, D. AIMD Evidence of Inner Sphere Adsorption of Glycine on a Stepped (101) Boehmite AlOOH Surface. *J. Phys. Chem. C* **2012**, *116*, 23418−23427.
(85) Ngouana-Wakou, B. F.; Cornette, P.; Corral Valero, M.; Costa, D.; Raybaud, P. An Atomistic Description of the γ-Alumina/Water Interface Revealed by Ab Initio Molecular Dynamics. *J. Phys. Chem. C* **2017**, *121*, 10351−10363.
(86) Zhang, Q.; Çağın, T.; van Duin, A.; Goddard, W. A.; Qi, Y.; Hector, L. G. Adhesion and nonwetting-wetting transition in the Al/α-Al_2O_3 interface. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2004**, *69*, 045423.
(87) Russo, M. F.; Li, R.; Mench, M.; van Duin, A. C. Molecular dynamic simulation of aluminum−water reactions using the ReaxFF reactive force field. *Int. J. Hydrogen Energy* **2011**, *36*, 5828−5835.
(88) Joshi, K. L.; van Duin, A. C. T. Molecular Dynamics Study on the Influence of Additives on the High-Temperature Structural and Acidic Properties of ZSM-5 Zeolite. *Energy Fuels* **2013**, *27*, 4481−4488.
(89) Sen, F. G.; Alpas, A. T.; van Duin, A. C. T.; Qi, Y. Oxidation-assisted ductility of aluminium nanowires. *Nat. Commun.* **2014**, *5*, 3959.
(90) Joshi, K. L.; Psofogiannakis, G.; van Duin, A. C. T.; Raman, S. Reactive molecular simulations of protonation of water clusters and depletion of acidity in H-ZSM-5 zeolite. *Phys. Chem. Chem. Phys.* **2014**, *16*, 18433−18441.
(91) Hong, S.; van Duin, A. C. Molecular Dynamics Simulations of the Oxidation of Aluminum Nanoparticles using the ReaxFF Reactive Force Field. *J. Phys. Chem. C* **2015**, *119*, 17876−17886.
(92) Gunkelmann, N.; Bringa, E. M.; Rosandi, Y. Molecular Dynamics Simulations of Aluminum Foams under Tension: Influence of Oxidation. *J. Phys. Chem. C* **2018**, *122*, 26243−26250.
(93) Ramírez, M.; González, R. I.; Baltazar, S. E.; Rojas-Nunez, J.; Allende, S.; Valdivia, J. A.; Rogan, J.; Kiwi, M.; Valencia, F. J. Thermal stability of aluminum oxide nanoparticles: role of oxygen concentration. *Inorg. Chem. Front.* **2019**, *6*, 1701−1706.
(94) Rosandi, Y.; Luu, H.-T.; Urbassek, H. M.; Gunkelmann, N. Molecular dynamics simulations of the mechanical behavior of alumina coated aluminum nanowires under tension and compression. *RSC Adv.* **2020**, *10*, 14353−14359.
(95) Chiche, D.; Chizallet, C.; Durupthy, O.; Chanéac, C.; Revel, R.; Raybaud, P.; Jolivet, J.-P. Growth of boehmite particles in the presence of xylitol: morphology oriented by the nest effect of hydrogen bonding. *Phys. Chem. Chem. Phys.* **2009**, *11*, 11310−11323.
(96) Mortier, W. J.; Ghosh, S. K.; Shankar, S. Electronegativity-equalization method for the calculation of atomic charges in molecules. *J. Am. Chem. Soc.* **1986**, *108*, 4315−4320.
(97) Verstraelen, T.; Ayers, P. W.; Van Speybroeck, V.; Waroquier, M. ACKS2: atom-condensed Kohn-Sham DFT approximated to second order. *J. Chem. Phys.* **2013**, *138*, 074108.
(98) ReaxFF force field format specification. https://www.scm.com/doc/ReaxFF/ffield_descrp.html (accessed on June 21, 2023).
(99) Chenoweth, K.; van Duin, A. C. T.; Goddard, W. A. ReaxFF Reactive Force Field for Molecular Dynamics Simulations of Hydrocarbon Oxidation. *J. Phys. Chem. A* **2008**, *112*, 1040−1053.
(100) The Amsterdam Modeling Suite. https://www.scm.com/amsterdam-modeling-suite (accessed on June 21, 2023).
(101) Plimpton, S. Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comput. Phys.* **1995**, *117*, 1−19.
(102) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; et al. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **2022**, *271*, 108171.
(103) ParAMS 2023.101 documentation. https://www.scm.com/doc/params (accessed on June 21, 2023).
(104) Gaus, M.; Cui, Q.; Elstner, M. DFTB3: Extension of the Self-Consistent-Charge Density-Functional Tight-Binding Method (SCC-DFTB). *J. Chem. Theory Comput.* **2011**, *7*, 931−948.
(105) Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; et al. The atomic simulation environment—a Python library for working with atoms. *J. Phys.: Condens. Matter* **2017**, *29*, 273002.
(106) Humphrey, W.; Dalke, A.; Schulten, K. VMD: Visual molecular dynamics. *J. Mol. Graph.* **1996**, *14*, 33−38.
(107) Kennes, K.; Kubarev, A.; Demaret, C.; Treps, L.; Delpoux, O.; Rivallan, M.; Guillon, E.; Méthivier, A.; de Bruin, T.; Gomez, A.; et al. Multiscale Visualization and Quantification of the Effect of Binders on the Acidity of Shaped Zeolites. *ACS Catal.* **2022**, *12*, 6794−6808.
(108) Perdew, J. P.; Ernzerhof, M.; Burke, K. Rationale for mixing exact exchange with density functional approximations. *J. Chem. Phys.* **1996**, *105*, 9982−9985.
(109) Hafner, J. Ab-initio simulations of materials using VASP: Density-functional theory and beyond. *J. Comput. Chem.* **2008**, *29*, 2044−2078.
(110) Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1994**, *50*, 17953−17979.
(111) Steinmann, S. N.; Corminboeuf, C. Comprehensive Benchmarking of a Density-Dependent Dispersion Correction. *J. Chem. Theory Comput.* **2011**, *7*, 3567−3577.
(112) LaBrosse, M. R.; Johnson, J. K.; van Duin, A. C. T. Development of a Transferable Reactive Force Field for Cobalt. *J. Phys. Chem. A* **2010**, *114*, 5855−5861.
(113) Iype, E.; Hutter, M.; Jansen, A. P. J.; Nedea, S. V.; Rindt, C. C. M. Parameterization of a reactive force field using a Monte Carlo algorithm. *J. Comput. Chem.* **2013**, *34*, 1143−1154.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX

(114) Boes, J. R.; Groenenboom, M. C.; Keith, J. A.; Kitchin, J. R. Neural network and ReaxFF comparison for Au properties. *Int. J. Quantum Chem.* **2016**, *116*, 979−987.
(115) Kokotailo, G. T.; Lawton, S. L.; Olson, D. H.; Meier, W. M. Structure of synthetic zeolite ZSM-5. *Nature* **1978**, *272*, 437−438.
(116) van Duin, A. C. T.; Baas, J. M. A.; van de Graaf, B. Delft molecular mechanics: a new approach to hydrocarbon force fields. Inclusion of a geometry-dependent charge calculation. *J. Chem. Soc., Faraday Trans.* **1994**, *90*, 2881.
(117) Hansen, N. Towards a new evolutionary computation. In *Advances on Estimation of Distribution Algorithms*; Lozano, J. A., Larrañaga, P., Inza, I., Bengoetxea, E., Eds.; Springer: Berlin, Heidelberg, 2006; pp 75−102.
(118) Hansen, N.; Müller, S. D.; Koumoutsakos, P. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evol. Comput.* **2003**, *11*, 1−18.
(119) Hansen, N. The CMA Evolution Strategy: A Tutorial. *arXiv* **2023**, arXiv:1604.00772.
(120) Raybaud, P.; Chizallet, C.; Pigeon, T. NOMAD dataset: Gamma-Alumina surfaces, 2021. https://nomad-lab.eu/prod/v1/gui/dataset/id/AlKosdljSIOsXeGXGt0TGg (accessed on November 16, 2022).
(121) Giagkiozis, I.; Fleming, P. Methods for multi-objective optimization: An analysis. *Inf. Sci.* **2015**, *293*, 338−350.
(122) Tian, Y.; Si, L.; Zhang, X.; Cheng, R.; He, C.; Tan, K. C.; Jin, Y. Evolutionary Large-Scale Multi-Objective Optimization: A Survey. *ACM Comput. Surv.* **2022**, *54*, 1−34.
(123) Blanchard, P.; Higham, D. J.; Higham, N. J. Accurately computing the log-sum-exp and softmax functions. *IMA J. Numer. Anal.* **2021**, *41*, 2311−2330.
(124) Palaz, D.; Synnaeve, G.; Collobert, R. Jointly Learning to Locate and Classify Words Using Convolutional Networks. *Proc. Interspeech* **2016**, *2016*, 2741−2745.
(125) Nian, R.; Liu, J.; Huang, B. A review On reinforcement learning: Introduction and applications in industrial process control. *Comput. Chem. Eng.* **2020**, *139*, 106886.
(126) Sutton, R. S.; Barto, A. G. Chapter Gradient Bandit Algorithms. In *Reinforcement Learning: An Introduction*; MIT Press, 2018; pp 37−40.
(127) Brown, K. S.; Sethna, J. P. Statistical mechanical approaches to models with many poorly known parameters. *Phys. Rev. E: Stat., Nonlinear, Soft Matter Phys.* **2003**, *68*, 021904.
(128) Waterfall, J. J.; Casey, F. P.; Gutenkunst, R. N.; Brown, K. S.; Myers, C. R.; Brouwer, P. W.; Elser, V.; Sethna, J. P. Sloppy-Model Universality Class and the Vandermonde Matrix. *Phys. Rev. Lett.* **2006**, *97*, 150601.
(129) Gutenkunst, R. N.; Waterfall, J. J.; Casey, F. P.; Brown, K. S.; Myers, C. R.; Sethna, J. P. Universally Sloppy Parameter Sensitivities in Systems Biology Models. *PLoS Comput. Biol.* **2007**, *3*, No. e189.
(130) Transtrum, M. K.; Machta, B. B.; Brown, K. S.; Daniels, B. C.; Myers, C. R.; Sethna, J. P. Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *J. Chem. Phys.* **2015**, *143*, 010901.
(131) Verstraelen, T.; Bultinck, P.; Van Speybroeck, V.; Ayers, P. W.; Van Neck, D.; Waroquier, M. The Significance of Parameters in Charge Equilibration Models. *J. Chem. Theory Comput.* **2011**, *7*, 1750−1764.
(132) Rahaman, O.; van Duin, A. C. T.; Bryantsev, V. S.; Mueller, J. E.; Solares, S. D.; Goddard, W. A. I.; Doren, D. J. Development of a ReaxFF Reactive Force Field for Aqueous Chloride and Copper Chloride. *J. Phys. Chem. A* **2010**, *114*, 3556−3568.
(133) Goken, E. G.; Joshi, K. L.; Russo, M. F.; van Duin, A. C. T.; Castleman, A. W. Effect of Formic Acid Addition on Water Cluster Stability and Structure. *J. Phys. Chem. A* **2011**, *115*, 4657−4664.
(134) Hu, X.; Schuster, J.; Schulz, S. E. Multiparameter and Parallel Optimization of ReaxFF Reactive Force Field for Modeling the Atomic Layer Deposition of Copper. *J. Phys. Chem. C* **2017**, *121*, 28077−28089.
(135) Bertels, L. W.; Newcomb, L. B.; Alaghemandi, M.; Green, J. R.; Head-Gordon, M. Benchmarking the Performance of the ReaxFF Reactive Force Field on Hydrogen Combustion Systems. *J. Phys. Chem. A* **2020**, *124*, 5631−5645.
(136) Daksha, C. M.; Yeon, J.; Chowdhury, S. C.; Gillespie, J. W. Automated ReaxFF parametrization using machine learning. *Comput. Mater. Sci.* **2021**, *187*, 110107.
(137) Yeon, J.; Chowdhury, S. C.; Daksha, C. M.; Gillespie, J. W. Development of Mg/Al/Si/O ReaxFF Parameters for Magnesium Aluminosilicate Glass Using an Artificial Neural Network-Assisted Genetic Algorithm. *J. Phys. Chem. C* **2021**, *125*, 18380−18394.
(138) Evans, D. J.; Holian, B. L. The Nose−Hoover thermostat. *J. Chem. Phys.* **1985**, *83*, 4069−4074.

https://doi.org/10.1021/acs.jctc.3c01009
J. Chem. Theory Comput. XXXX, XXX, XXX−XXX