


# Machine Learning Force Fields: Construction, Validation, and Outlook

V. Botu,<sup>\*,†</sup> R. Batra,<sup>‡</sup> J. Chapman,<sup>§</sup> and R. Ramprasad<sup>‡</sup>

<sup>†</sup>Department of Chemical & Biomolecular Engineering, <sup>‡</sup>Department of Materials Science & Engineering, and <sup>§</sup>Department of Materials Science, University of Connecticut, Storrs, Connecticut 06269, United States

## Supporting Information

**ABSTRACT:** Force fields developed with machine learning methods in tandem with quantum mechanics are beginning to find merit, given their (i) low cost, (ii) accuracy, and (iii) versatility. Recently, we proposed one such approach, wherein, the vectorial force on an atom is computed directly from its environment. Here, we discuss the multistep workflow required for their construction, which begins with generating diverse reference atomic environments and force data, choosing a numerical representation for the atomic environments, down selecting a representative training set, and lastly the learning method itself, for the case of AI. The constructed force field is then validated by simulating complex materials phenomena such as surface melting and stress–strain behavior, that truly go beyond the realm of *ab initio* methods, both in length and time scales. To make such force fields truly versatile an attempt to estimate the uncertainty in force predictions is put forth, allowing one to identify areas of poor performance and paving the way for their continual improvement.



## INTRODUCTION

Materials modeling approaches largely fall in two broad categories: one based on quantum mechanical methods (e.g., density functional theory), and the other based on semi-empirical analytical interatomic potentials or force fields (e.g., Stillinger-Weber potentials, embedded atom method, etc.).<sup>1–7</sup> Choosing between the two approaches depends on which side of the cost–accuracy trade-off one wishes to be at. Quantum mechanical methods (also referred to as *ab initio* or first-principles methods) are versatile and offer the capability to accurately model a range of chemistries and chemical environments. However, such methods remain computationally very demanding. Practical and routine applications of these methods at the present time are limited to studies of phenomena whose typical length and time scales are of the order of nanometers and picoseconds, respectively. Semi-empirical methods capture the essence of these interatomic interactions in a simple manner via parametrized analytical functional forms, and thus offer inexpensive solutions to the materials simulation problem. However, their applicability is severely restricted to the domain of chemistries and chemical environments intended, or considered during parametrization.<sup>8</sup> It is unclear whether the underlying framework allows for a systematic and continuous improvement in the predictive capability of newer environments.

The present contribution pertains to a data-driven approach by which flexible and adaptive force fields may be developed, potentially addressing the challenges posed. By using carefully created benchmark data (say, from quantum mechanics based materials simulations) as the starting point, nonlinear

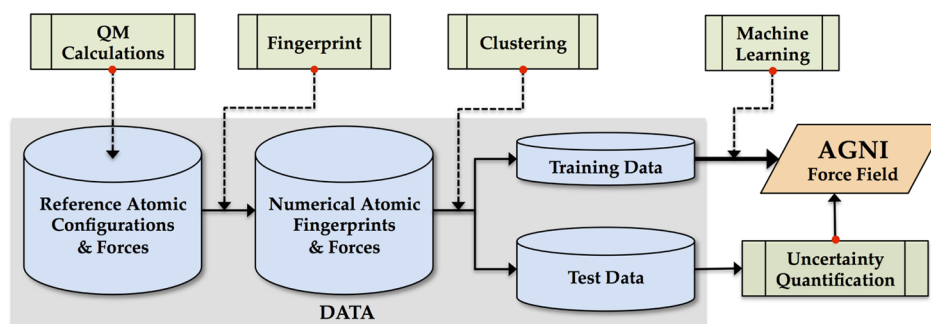
associations between atomic configurations and potential energies (or forces, more pertinent to the present contribution) may be learned by induction.<sup>9–11</sup> This data-driven paradigm, popularly referred to as machine learning, has been shown by many groups to lead to viable pathways for the creation of interatomic potentials that (1) surpass conventional interatomic potentials both in accuracy and versatility, (2) surpass quantum mechanical methods in cost (by orders of magnitude), and (3) rival quantum mechanics in accuracy,<sup>12–15</sup> at least within the configurational and chemical domains encompassed by the benchmark data set used in the training of the potential.

A new recent development within the topic of machine learning based interatomic potentials is the realization that the vectorial force experienced by a particular atom may be learned and predicted directly given just a configuration of atoms.<sup>16–18</sup> This capability is particularly appealing as the atomic force is a local quantity purely determined by its immediate environment, in contrast to conventional methods that predict the total potential energy, which is a global property of the system as a whole. A large body of materials simulations, such as geometry optimization and molecular dynamics simulations, require the atomic force as the sole necessary input ingredient.<sup>1</sup> Note that partitioning the total potential energy into individual atomic contributions, conventionally adopted in semiempirical interatomic potentials, is a matter of convenience of construction, rather than being a fundamental requirement.

**Received:** October 30, 2016

**Revised:** December 3, 2016

**Published:** December 7, 2016



**Figure 1.** Flowchart illustrating the workflow in constructing AGNI force fields: generating reference atomic configurations and forces with quantum mechanical methods, fingerprinting atomic environments, rational selection of training and test data sets, mapping atomic fingerprints to forces using machine learning methods, and quantifying uncertainty in predictions made.

This article deals specifically with using machine learning methods to create an atomic force prediction capability, i.e., a force field. As recently pointed out, this force field is adaptive (i.e., new configurational environments can be systematically added to improve the versatility of the force field, as required), generalizable (i.e., the scheme can be extended to any collection of elements for which reliable reference calculations can be performed), and neighborhood informed (i.e., a numerical fingerprint that represents the atomic environment around the reference atom is mapped to the atomic force with chemical accuracy).<sup>16,17</sup> The force field is henceforth dubbed AGNI.

The workflow in constructing AGNI force fields includes five key steps. These include: (1) creation of a reference data set derived from a plethora of diverse atomic environments of interest, and the corresponding atomic forces computed using a chosen quantum mechanical method, (2) fingerprinting every atomic environment, in a manner that will allow the fingerprint to be mapped to atomic force components, (3) choosing a subset of the reference data set (the “training” set), using clustering techniques to optimize the learning process while ensuring that the training set represents the diversity encompassed by the original reference data set, (4) learning from the training set, thus leading to a nonlinear mapping between the training set fingerprints and the forces, followed by testing the learned model on the remainder of the data set using best-statistical practices, and (5) finally, estimation of the expected levels of uncertainty of each force prediction, so that one may determine when the force field is being used outside its domain of applicability. The entire workflow involved in the construction of AGNI force fields is portrayed schematically in Figure 1.

In our previous work, a preliminary version of the AGNI force field for Al was used to demonstrate its capability with respect to predicting structural, transport or vibrational properties of materials.<sup>17</sup> Here, we further extend its scope, by including more diverse atomic environments and its ability to simulate even more complex phenomena, such as surface melting and stress–strain behavior. Furthermore, although AGNI is built to provide atomic forces, we demonstrate that accurate total potential energies can be retrieved either during the course of a molecular dynamics simulation or along a reaction coordinate, through appropriate integration of atomic forces.


Additional comments pertaining to the last step of the workflow in Figure 1 are in order. Uncertainty quantification is essential to recognize when the force field is operating outside its domain of applicability. Ideally, the larger the uncertainty of

the force prediction for an atom in a given environment, the greater is the likelihood that the environment is “new”. By imposing a threshold and monitoring the uncertainty we may wish to augment the training set with the corresponding new atomic environment(s) and follow the workflow in Figure 1. This helps build force fields that are truly adaptable. Initial steps toward quantifying this uncertainty in force predictions are undertaken.

The rest of the paper is organized as follows. In the first half of the work we guide the readers through the rigors of each step in the force field construction workflow, shown in Figure 1, to develop a general-purpose AI force field. The applicability of the force field is then validated by demonstrating its use in atomistic simulations. A discussion on measures to estimate uncertainties in force predictions made is then put forth. Lastly, we conclude with an outlook on using machine learning force fields in the field of atomistic materials modeling, and the challenges that yet remain to be addressed.

## ■ COMPUTATIONAL DETAILS

**Generating Reference Data.** The construction of AGNI force fields begins with data. We start by building several periodical and nonperiodical equilibrium configurations (c.f., Figure 2), such as (i) defect free bulk, (ii) surfaces, (iii) point defects, vacancies and adatoms, (iv) isolated clusters, (v) grain boundaries, (vi) lattice expansion and compression, and (vii) edge type dislocations. These configurations are so chosen to mimic the diverse environments an atom could exist in and forms a critical first step in constructing generalizable force fields. It is by no means a complete list and one could continuously add nonredundant configurations to it (methods to identify such redundancies are discussed later). The vectorial force components on each atom, in the equilibrium configurations amassed, are then computed by quantum mechanical based density functional theory (DFT) calculations.<sup>19,20</sup> To correctly describe the nonequilibrium behavior of an atom, in response to a perturbation due to thermal vibrations, pressure or other sources, it is equally necessary to construct nonequilibrium atomic environments, to learn the complete array of forces experienced by an atom. A convenient and quick means to sampling such nonequilibrium environments is with *ab initio* molecular dynamics (MD) simulations.<sup>21</sup> Here, starting with the equilibrium configurations, MD simulations in the micro-canonical ensemble were then carried out with multiple initial velocity conditions (between 200–1250 K). One could equally use multiple constant temperature MD simulations as well, to source the reference



**Figure 2.** Reference configurations used to sample atomic environments for training and testing of AGNI force fields; (i) bulk, (ii) surfaces, (iii) defects (vacancies and adatoms), (iv) isolated clusters, (v) grain boundaries, (vi) lattice expansion and compression, and (vii) dislocation.

data. For large reference configurations, where MD simulations are restricted, initially the atoms are randomly perturbed to coerce the dynamics into quickly sampling nonequilibrium environments. The combination of *ab initio* MD and random perturbations resulted in a diverse set of reference atomic environments and forces (cf., Table 1), needed to learn (indirectly) the underlying potential energy surface.

From within the millions of reference atomic environments collected, a subset of them are chosen as training environments to construct the force fields. The particular choice of environments plays a critical role in the generalizability of

**Table 1. Atomic Environment Makeup for the Five Datasets; A, B, C, D, and E<sup>a</sup>**


data set	atomic envs from reference configurations	no. of envs
A	Defect free bulk fcc and bcc.	20385
B	Data set A + (100), (110), (111), (200), and (333) surfaces.	211255
C	Data set B + Defects in bulk fcc with 1, 2, and 6 randomly distributed vacancies and adatom on (100), (110) and (111) surfaces.	1502856
D	Data set C + Isolated clusters of 5 Å, 8 Å, 10 Å, and 12 Å.	586679
E	$\Sigma 3$ (111), $\Sigma 5$ (210), $\Sigma 5$ (310), $\Sigma 13$ (320), and $\Sigma 13$ (510) grain boundaries, varying lattice vectors by $\pm 7\%$ of equilibrium, edge dislocation along (112) direction.	394116

<sup>a</sup>For each dataset we generate a training and test set (except for dataset E, where only a test set is created), the former used to construct the force field and the later to validate it. The number of new environments added is indicated in the last column.

such data-driven force fields. To better understand such limits imposed by data choices, we construct four data sets, labeled as A, B, C, and D, with increasing complexity and diversity of atomic environments contained (cf., Table 1). For each data set, training and test sets were created, the former used to construct the force field and the later used to validate its predictive prowess. Also a fifth data set, E, consisting of configurations never used during force field construction (see Table 1) was created, solely to demonstrate the transferability of AGNI force fields.

All force and MD calculations were done using VASP, a plane-wave based DFT software.<sup>22,23</sup> The PBE functional to treat the electronic exchange-correlation interaction, the projector augmented wave potentials, and plane-wave basis functions up to a kinetic energy cutoff of 520 eV were used.<sup>24,25</sup> A  $14 \times 14 \times 14$   $\Gamma$ -centered k-point mesh was used for the primitive Al unit cell and scaled according to the unit cell size. A time step of 0.5 fs was chosen for the MD simulations. Methods to access the reference DFT data files used in this work are discussed in the Supporting Information.

**Fingerprinting Reference Environments.** Choosing a representation for an atom and its environment is the most critical step in the entire workflow. In order to learn the vectorial force components,  $F^u$ , where  $u$  refers to any arbitrary direction, necessitates a numerical representation that conforms with this directional dependence. Further, it should also remain invariant to the basic atomic transformation operations, such as translation, rotation, or permutation. One such representation (commonly referred to as a fingerprint) with the necessary prerequisites is given below



**Figure 3.** Schematic demonstrating the scalar projection for an atom  $i$  (the reference atom) and one of its neighbor (1) along a direction  $u$ . To generate the final fingerprint for atom  $i$ , a summation over the atoms within the cutoff sphere, as indicated by the dashed line, is considered.

$$V_i^u(\eta) = \sum_{j \neq i} \frac{r_{ij}^u}{r_{ij}} \cdot e^{-(r_{ij}/\eta)^2} \cdot f_d(r_{ij}) \quad (1)$$

Here,  $r_{ij}$  is the distance between atoms  $i$  and  $j$  ( $\|\mathbf{r}_j - \mathbf{r}_i\|$ ), while  $r_{ij}^u$  is a scalar projection of this distance along a direction  $u$  (c.f., Figure 3).  $\eta$  is the Gaussian function width.

$f_d(r_{ij}) = 0.5 \left[ \cos\left(\frac{\pi r_{ij}}{R_c}\right) + 1 \right]$  is a damping function for atoms within the cutoff distance ( $R_c$ ) and is zero elsewhere. The summation in eq 1 runs over all neighboring atoms within an arbitrarily large  $R_c$  (8 Å, in the present work).

To better understand the fingerprint described in eq 1, one can deconvolute it into three subcomponents (as separated by “.”). The exponential term ( $e^{-(\frac{r_{ij}}{\eta})^2}$ ) imposes a coordination shell around an atom  $i$  with  $\eta$  describes the extent of the shell. By using multiple such  $\eta$  values both nearby and distant coordination information are contained within the fingerprint. In this work,  $\eta$ 's were sampled on a logarithmic grid between [0.8 Å, 16 Å], ensuring a sufficient description of the neighbor interactions. One could also use the peak positions of a radial distribution function as a starting point in choosing the  $\eta$  values. The normalized scalar projection term ( $\frac{r_{ij}^u}{r_{ij}}$ ) adds directionality to the fingerprint by selectively resolving the coordination information along the desired direction ( $u$ ), and is necessary to map the individual force components. Lastly, the damping function ( $f_d(r_{ij})$ ) diminishes the influence of far away atoms smoothly. The combination of these three features makes this particular choice of representation suitable for mapping atomic force components. Similar coordination based fingerprints have been developed in the past.<sup>12,26</sup> However, these were tailored for the purpose of mapping the total potential energy (a scalar quantity) for a given configuration of atoms, unlike the vectorial force components as done here.

Further, the atomic fingerprint representation chosen conforms with the required invariance operations, such as permutation, translation, and rotation of atoms. For instance consider a reference atom,  $i$ , and its neighboring atoms within a cutoff sphere (as shown in Figure 3). Information pertaining to atoms neighboring atom  $i$  are passed into the summand of eq 1 as pairwise distances, thereby, permutation or translation of atoms does not alter  $V_i^u$ . In the case of rotations, both the fingerprint and vectorial force components change in a manner governed by the rotation matrix. For example, by rotating atoms along the  $z$ -axis one can redefine the forces (shown here

for the force but equally applicable to the fingerprint) along the Cartesian directions as

$$\begin{bmatrix} F^{x'} \\ F^{y'} \\ F^{z'} \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} F^x \\ F^y \\ F^z \end{bmatrix} \quad (2)$$

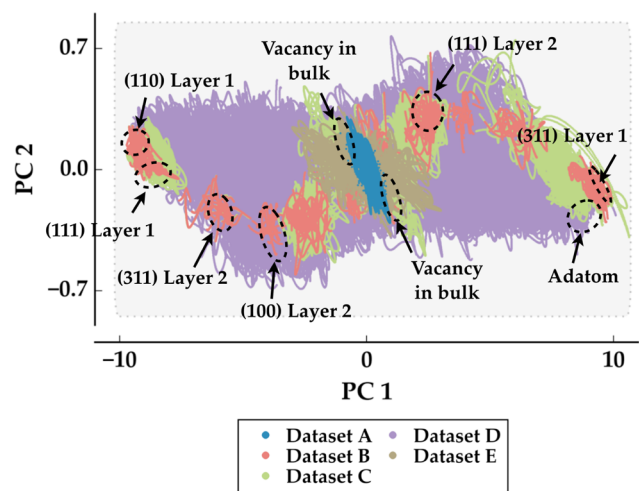
Nevertheless, the magnitude of the net force before and after rotation remains the same, as one would expect. The concurrence to this rotational behavior implies that both the forces and fingerprints transform in an identical manner upon rotation.

An added advantage of this rotational property is that one can expand upon the pre-existing reference database beyond the three vectorial components (i.e., the Cartesian direction) obtained from DFT methods. By defining a grid of  $\theta$  values in eq 2, arbitrary projections of the atomic fingerprint and forces can be computed without any additional costly ab initio calculations. Similarly, multiple values for the azimuth angle,  $\phi$  (taken as 0 in eq 2), can be chosen, adding to the diversity and completeness of possible reference atomic environments and forces. However, this builds in extensive redundancies. Training a force field on millions of reference atomic environments is impractical, computationally very demanding, and might lead to misbehaved models, therefore, further down-sampling from within this big pool of data is an essential step in the construction workflow, as discussed shortly.

Another important aspect of the fingerprint is that it remains unique for the diverse atomic environment situations. Numerically this implies that identical fingerprints should map to the same atomic force value. Here, we do this by using multiple  $\eta$  values. In the limit that number of  $\eta$  values tends to  $\infty$  uniqueness can be ensured, nevertheless, for all practical purposes one can make do with a much smaller subset, as determined by running convergence tests.

**Clustering Reference Data.** The next step in the construction workflow is to select a representative set of atomic environments for training purposes. To do so, it is necessary to identify the redundant and noncontributing data points from within the millions sampled. An obvious place to start is by comparing among the individual atomic fingerprints. However, given its high-dimensionality understanding or unraveling the fingerprint directly is nontrivial. Therefore, we rely on dimensionality reduction techniques such as principal component analysis (PCA) to project  $V_i^u$  onto a lower dimension space.<sup>27</sup>

In PCA the original atomic fingerprint is linearly transformed into uncorrelated and orthogonal pseudo variables, also known as principal components (PCs). Often times the information content contained within the original fingerprint can be captured by a few such PCs. To demonstrate this, here, for all the reference atomic environments we compute and transform an 8-dimensional fingerprint (the rationale for which shall be discussed shortly). Two such PCs captured more than 99% of the information content of the original fingerprint, allowing us to visualize the atomic environments on a two-dimensional manifold known as a scores plot (cf., Figure 4). Immediately, we observe clustering of the atomic finger-



**Figure 4.** Projection of an 8 dimensional (i.e., 8  $\eta$  values) atomic fingerprint for all entries within data sets A, B, C, D, and E onto a 2-dimensional manifold, transformed by principal component analysis. Identical atomic environments cluster together and highlights the extent of redundancies within the reference data set.

prints that correspond to similar neighborhood environments. For clarity environments corresponding to a few such cases, e.g., adatoms, surfaces, vacancies, etc., are labeled. Further, by color coding atoms according to the data set they were sampled from, i.e., A, B, C, D, or E, we qualitatively observe their extent of diversity. For instance, data set D (cf., Figure 4) spans a diverse set of atomic environments as it populates majority of the space, suggesting that isolated cluster configurations are a good starting point to sample reference data from. Interestingly, atomic environments from data set E lie within the domain of data set D. This suggests that a force field trained on data set D should accurately predict the forces for environments in data set E, though they were never explicitly included during training.

Thus, far, the dimensionality reductions methods used provided a visual means to identifying redundancies, but for an efficient force field construction an automated sampling of the nonredundant training environments is necessary. One way of doing so is to choose the data randomly. Unfortunately, this biases sampling according to the underlying probability distribution of the data set and fails to sample sparsely populated regions. To avoid such irregularities, here we adopt a simple grid-based sampling on the PCA space. The PCA transformed data is split into uniform subgrids, the bounds of which are determined by the minimum and maximum of the relevant PCs (which can be more than 2). Training points are then randomly sampled from within each subgrid. By using a

fine grid one can ensure uniform and diverse sampling from all regions of the PC space. Note that in the limit the grid size becomes large this approach is equivalent to the random sampling approach. Also, the test sets for validation purposes are generated in a similar manner from the remaining nonsampled data. Other dimensionality reduction algorithms, such as kernel-PCA<sup>28</sup> or multidimensional scaling,<sup>29</sup> could similarly be adopted to sample for a representative data set. Irrespective of the choice, such clustering methods are critical as the learning and prediction cost scales as  $O(n^3)$  and  $O(n)$ , respectively (where  $n$  is the training data set size).

**Learning Algorithm.** The next vital ingredient required in putting together a predictive framework is the learning algorithm itself. Deep learning neural networks<sup>30</sup> and nonlinear regression processes<sup>13</sup> have been the methods of choice for models describing atomic interactions. Their capability to handle highly nonlinear relations, as is in the case of mapping an atom's environment to the force it experiences, makes them a suitable choice here as well. Here, we choose nonlinear kernel ridge regression (KRR) method as the machine learning workhorse.<sup>11,31</sup> KRR works on the principle of (dis)similarity, wherein, by comparing an atom's fingerprint ( $V_i^u(\eta)$ ) with a set of reference cases, an interpolative prediction of the  $u$ th component of the force ( $F_i^u$ ) can be made

$$F_i^u = \sum_t \alpha_t \cdot \exp \left[ -\frac{(d_{i,t}^u)^2}{2l^2} \right] \quad (3)$$

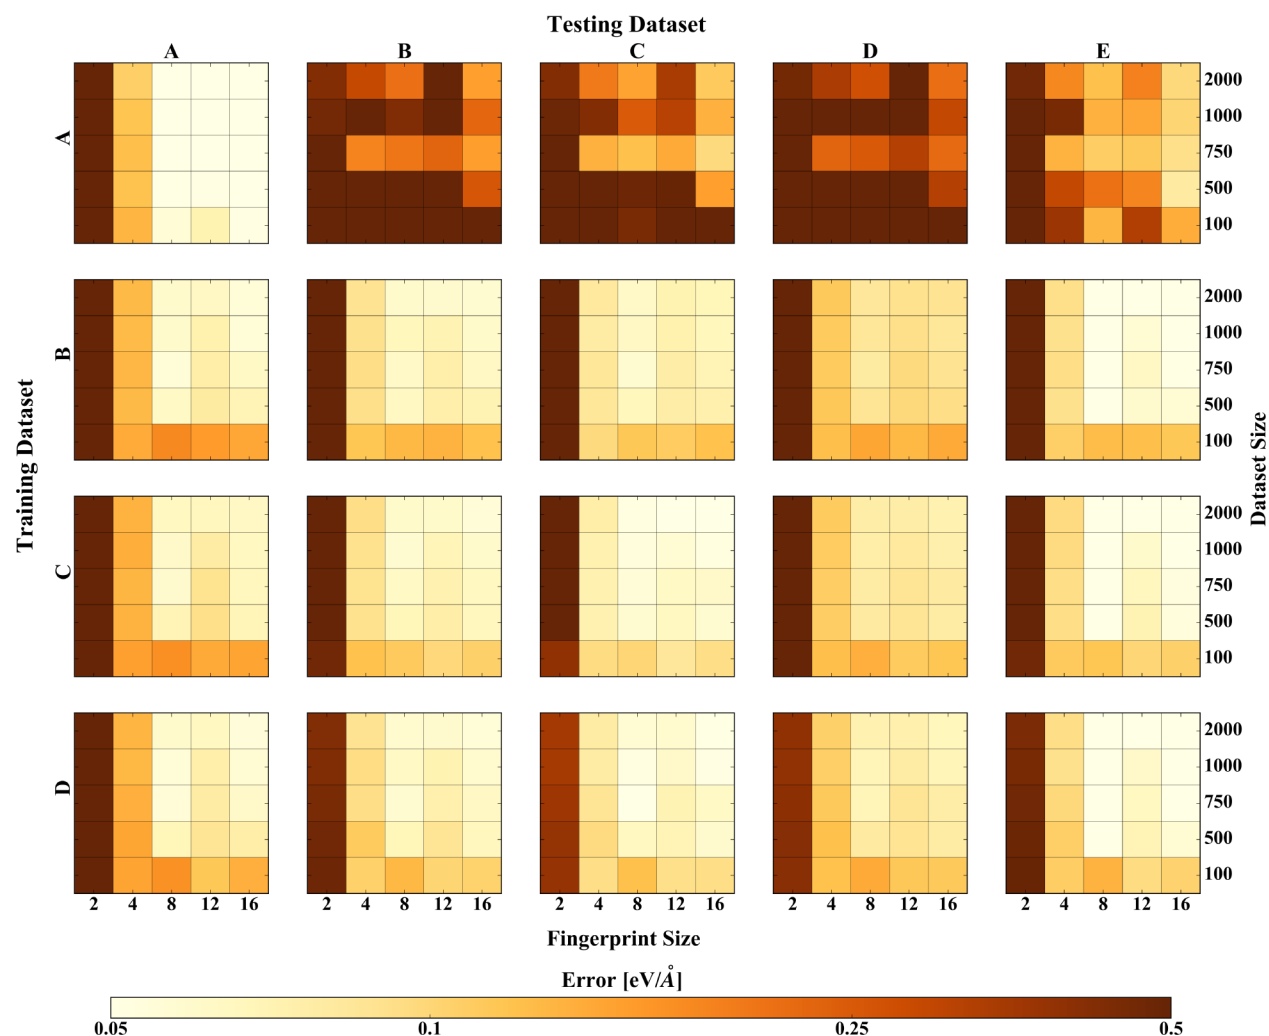
Here,  $t$  labels each reference atomic environment, and  $V_t^u(\eta)$  is its corresponding fingerprint.  $N_t$  is the total number of reference environments considered.  $d_{i,t}^u = \|V_i^u(\eta) - V_t^u(\eta)\|$  is the Euclidean distance between the two atomic fingerprints, though other distance metrics can be used.  $\alpha_t$ s and  $l$  are the weight coefficients and length scale parameter, respectively. The optimal values for  $\alpha_t$ s and  $l$  are determined during the training phase, with the help of cross-validation and regularization methods. For further details concerning the learning algorithm the reader is directed to these sources.<sup>10,32,33</sup>

Finally, in order to evaluate the performance of a developed force field, three error metrics—mean absolute error (MAE), maximum absolute error (MAX), and the standard deviation (in particular  $2\sigma$ )—were chosen. Relying on multiple metrics reduces any bias, unknowingly, introduced during model selection as shall be discussed shortly.

## RESULTS AND DISCUSSION

**Constructing the Force Field.** At this stage all the pieces required to construct AGNI force fields, as illustrated by the flowchart in Figure 1, have been laid out. In the upcoming sections we discuss how one chooses the appropriate number of  $\eta$  values to adequately describe an atomic environment, and the minimum number of training environments to use, as needed to construct accurate force fields. The accepted force field is then put to the test by predicting forces on atoms outside the domain of training environments used, to ensure its generalizability.

**Convergence Tests.** The first step to attaining an optimal force field is to ensure convergence with respect two parameters: (i) the number of  $\eta$  values used for the atomic fingerprint and (ii) the training data set size. As mentioned earlier the number of  $\eta$  values governs the resolution with which an atom's local coordination environment is described,




**Figure 5.** Heat maps illustrating model error (mean absolute error) as a function of number of  $\eta$  values and the training data set size. The fingerprint was varied from 2 to 16  $\eta$  values, while the training data set size was varied from 100 to 2000 environments. We report the MAE for models trained on data sets A, B, C, and D and tested on data sets A, B, C, D, and E. For example the top row corresponds to models trained on data set A, while each column corresponds to a test data sets of the five cases. The errors quickly converge for a fingerprint with 8  $\eta$  values and a training size of 1000 diverse environments.

while the size (and choice, as shall be elaborated in the next section) of training data governs AGNI's interpolative predictive capability. In order to identify this optimal parameter set, we systematically increase the fingerprint resolution from 2 to 16  $\eta$  values and the training data set size from 100 to 2000 atomic environments, while monitoring test set error (in this case the chosen error metric was the MAE). To remind the reader,  $\eta$  values were sampled on a logarithmic grid between [0.8 Å, 16 Å], while training data environments were sampled using the PCA projection followed by a grid-based sampling.


For each of the four training data sets (A, B, C, and D), partitioned in the data generation stage, and for all combinations of the convergence parameters an AGNI force field was constructed. Each force field is then validated on the respective test data sets (A, B, C, D, and E). The force fields are denoted as  $M_{ij}^k$ , where  $i$  and  $j$  label the training and test environments used, respectively (the superscript is omitted when referring to the training environments only). Figure 5 illustrates heat maps of the error for different training and test data sets, and convergence parameter combinations. Two key findings stand out: (i) by increasing the fingerprint resolution the error drops and quickly converges below  $\sim 0.05$  eV/Å

(expected chemical accuracy) and (ii) increasing the training data set size reduces error only beyond a reasonable fingerprint resolution. For example, in  $M_{AC}^C$  increasing the training data set size for a fingerprint with 2 or 4  $\eta$  values has no effect on the predictive capability. Such a manifestation implies that 8 or more  $\eta$  values are required to “uniquely” discern among the atomic environments, in order for the learning algorithm to work. Nevertheless, this relation only holds for force fields used in an interpolative manner, as seen in the failure of  $M_{AV}^B$ ,  $M_{AV}^C$ ,  $M_{AV}^D$ , or  $M_{AV}^E$ . Here, the diversity in the training data chosen plays a more prominent role in governing performance, as shall be elaborated in the next section. Overall, we find that a fingerprint of 8  $\eta$  values and a training size of 1000 atomic environments is sufficient, beyond which the models exhibit diminishing returns, i.e., increased model training costs with no significant drop in model error, and are the parameters chosen for all subsequent discussions. The computational burden of each AGNI prediction is  $\sim 0.1$  ms/atom/core, while DFT costs  $\sim 1$  ks/atom/core.

**Training Data Choice.** The particular choice of atomic environments included during training is a crucial factor, as briefly alluded to earlier. Given that the learning algorithm is



**Figure 6.** (a) Mean absolute error, (b) maximum absolute error, and (c)  $2\sigma$  standard deviation error metric, for models trained on A, B, C, and D and tested on data set A, B, C, D, and E. The errors are reported for an 8 dimensional fingerprint vector (i.e., 8  $\eta$  values) chosen to represent the environment around an atom, and a training set size of 1000 environments obtained from the PCA grid-based sampling.



**Figure 7.** Parity plots comparing force predictions, for test environments in data set E, made with the AGNI model,  $M_D$ , EAM interatomic potential, and DFT. The plots have been further separated according to the configuration they were sampled from, (a) grain boundaries, (b) lattice expansion/compression, and (c) edge dislocation, respectively.


interpolative by nature, a force field trained say only on bulk type environments ( $M_A$ ) cannot predict the forces corresponding to other environments types, e.g., data sets with surfaces and other features— $M_A^B$ ,  $M_A^C$ ,  $M_A^D$ , or  $M_A^E$ . By increasing the diversity in training environments,  $M_B$ ,  $M_C$ , and  $M_D$ , we make the force fields more generalizable once the optimal parameters are chosen, as given by their low test error in Figure 5. Surprisingly, it appears as though predictions made with  $M_B$  are equally as good as  $M_C$  or  $M_D$ . However, this is purely a manifestation of using the MAE as the error metric. Along with the MAE, we report test set errors computed with two other metrics: MAX and  $2\sigma$ , as illustrated in Figure 6 (shown only for the optimal 8-component fingerprint and 1000 training atomic environments). For  $M_B$ , with MAX as the metric, the prediction error is high outside its domain of applicability (test set C, D, or E), and a similar behavior is observed for  $M_C$ . It should be recognized that MAX reports the worst prediction made, while MAE reports a mean error skewed by test set size. By combining the two metrics with the actual variance in the errors, as measured by the  $2\sigma$  metric, we can ensure that the error is indeed under control. We observe that in  $M_D$ , by sampling atomic environments from a very diverse set of configurations, all the error metrics are low, and the force field is highly generalizable and is the force field used in subsequent discussions.

**Testing out of domain configurations.** The configurations contained in data set E, grain boundaries, lattice expansion and compression, and dislocations, were never “observed” during the training phase. Being able to accurately predict the forces will further demonstrate the fidelity in using a local-neighborhood based force predictive capability. The PCA scores plot, shown in Figure 4, provided a glimpse of what one could expect. Given that the transformed atomic fingerprints for data set E lie within the domain of environments from data set D, one can expect the force predictions made by  $M_D$  to be

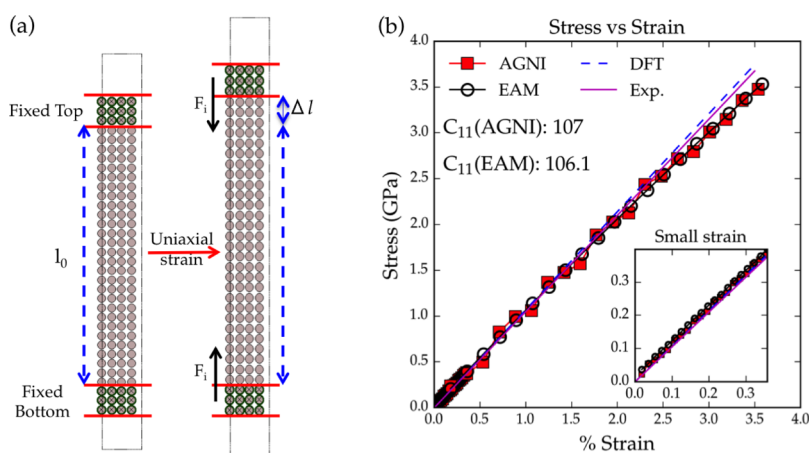
interpolative and thus accurate. However, a more stringent test is to predict forces on all the atoms in data set E and compare them to those obtained by DFT methods, as is done and shown in Figure 7. For all three cases, the AGNI predicted forces are in excellent agreement with DFT. This demonstrates the intended goal of AGNI force fields, i.e., to retain quantum mechanical accuracy, be computationally inexpensive, and remain generalizable. The last feature in particular, generalizability, is often lacking with traditional semiempirical methods. For comparison, we recompute the forces for atoms in data set E using traditional semiempirical potentials. Here, we particularly use an AI EAM potential,<sup>34</sup> as it accurately captures interactions in close-packed metallic type systems. As with AGNI force field, EAM methods equally predict forces accurately for grain boundaries and lattice expansion/compression but fails for dislocation type of environments.<sup>8</sup>

**Validating the Force Field.** Having demonstrated a robust scheme that allows accurate atomic force predictions for a diverse set of situations, in the subsequent sections, we demonstrate the true prowess of such AGNI force fields in facilitating atomistic simulations. Our previous work provided a glimpse of such simulations, whereby, structural optimization, vibrational property estimation, and simple MD simulations of materials were undertaken.<sup>17</sup> Here, we extend the realm of such force fields to simulate more complex atomistic phenomena, such as surface melting and stress–strain behavior. These are particularly challenging as the atoms traverse through a multitude of environments, and an accurate prediction of the forces requires undertaking the rigorous construction workflow discussed thus far.

The simulations were carried out using the LAMMPS molecular dynamics code.<sup>35</sup> The source code and force field files required to carry out the simulations are provided as supplemental files.



**Figure 8.** (a) Schematic of the (111) Al surface model and the regions classified as surface or bulk atoms. (b) The Lindemann index (LI) as a function of temperature simulated with both the AGNI and EAM force fields. Melting occurs once the LI transitions from a slow linear increase to a sudden spike. With the AGNI force field the surface begins to melt  $\sim 950$  K, and propagates to the bulk by  $\sim 1100$  K.



**Figure 9.** (a) Schematic of the surface model and procedure adopted to extract stress–strain behavior. (b) Stress vs strain behavior of uniaxially strained Al as computed using forces predicted by the AGNI and an EAM force field. The inset shows stress at low strain ( $<0.35\%$ ).

**Melting Behavior of an Al (111) Surface.** The melting temperature of condensed matter is a property often estimated by MD simulations. Here, starting with a surface model with over 1000 atoms and dimensions of  $19 \text{ \AA} \times 17 \text{ \AA}$  (cf., Figure 8a), constant temperature MD simulations were carried out for over 50 ps and across a temperature range of 300–1300 K, to estimate the melting temperature. In traditional MD simulations energy is used as the metric to distinguish between a solid and liquid state. Since this metric is not at our disposal we rely on the Lindemann Index (LI) order parameter instead. The LI measures the thermal perturbations of atoms: for solids this value is around  $\sim 0.03$  while for liquids it is  $\sim 0.13$ . A sudden increase in the LI as a function of temperature is attributed to a solid–liquid phase transition and can thus be utilized in MD simulations to estimate the melting temperature.<sup>36–38</sup> The LI is defined as


$$LI = \frac{1}{N} \sum_i \frac{1}{N-1} \sum_{j \neq i} \frac{\sqrt{\langle r_{ij}^2 \rangle - \langle r_{ij} \rangle^2}}{\langle r_{ij} \rangle} \quad (4)$$

where  $r_{ij}$  is the distance between atom  $i$  and  $j$ ,  $N$  is the total number of atoms, and  $\langle \dots \rangle$  indicate time averaged quantities.<sup>39</sup> For the range of temperatures considered, the LI was computed using eq 4 and reported in Figure 8b (red lines). Further, we distinguish between the surface and bulk LI values (cf., 8a) as it

is well-known that the melt front initiates at the surface and propagates inward. Starting at 300 K, the LI rises due to a systematic increase in thermal vibrations up to 900 K. Between 900–1000 K, the sudden increase in LI signifies the onset of surface melting, which then propagates into the bulk by 1200 K. With AGNI force fields this onset of melting is observed at  $\sim 950$  K, similar to the known experimental value at  $\sim 933$  K for Al.<sup>40</sup> Even though the training environments used in building the force field did not explicitly contain Al in a liquid state, by including high temperature MD reference data we were able to predict forces for environments in such extreme conditions. Similar LI curves computed by an EAM force field, as pure ab initio studies of this size and time scale are intractable, yielded an overestimated melting temperature of  $\sim 1100$  K.

**Stress–Strain Behavior of an Al (001) Surface.** Another important material property determined by atomistic simulations is the stress–strain behavior. Beginning with a (001) fcc Al surface, of dimensions  $8 \text{ \AA} \times 8 \text{ \AA} \times 80 \text{ \AA}$ , the surface atoms (fixed in their position) are displaced,  $\Delta l$ , resulting in a uniaxial strain along the surface normal [001] as illustrated in Figure 9a. Atoms within the strained region are then relaxed to minimize the forces acting upon them. This imposes a net force ( $F_s$ ) on the fixed surface atoms toward the bulk (cf., Figure 9a). Using the stress tensor relation,  $\sigma_{s,p} = \frac{F_s}{A_p}$ , for a plane  $p$  with an area





**Figure 10.** Evolution of energy of bulk fcc Al using the (a) AGNI with eq 6 and (b) EAM potential. DFT computed energies of a few configurations are also marked in (a) for validation. In all the cases, the energy is in reference to that of the perfect fcc Al. The energy is conserved with time and maintains the correct ordering with temperature.

$A_p$ , across varying values of strain,  $\epsilon_{s,p} = \frac{\Delta l}{l_0}$ , one can deduce the stress–strain behavior of the material. In Figure 9b we report the computed stress for varying strain deformations. The slope of this curve, for the direction considered, yields the  $C_{11}$  elastic coefficient, a property that can be compared with atomistic theories. Using the AGNI force field we report a  $C_{11}$  value of 107 GPa, which is in good agreement with a past ab initio result of 105 GPa.<sup>41</sup> Clearly, this suggests that besides forces the force field can predict their derivatives, i.e., the stresses, at quantum mechanical accuracy as well. Further, we recomputed the  $C_{11}$  value with an EAM potential, resulting in a value of 106 GPa. Though, the combination of AGNI force prediction and stress relation result in quantum mechanically accurate elastic coefficients, the procedure laid out can only describe the stress along nonperiodical directions, a limitation of the force based implementation.

**Energy Prediction.** We now briefly touch upon the topic of energy. Energy is a unique, and important, global quantity describing the state of a configuration of atoms. It is often used to ensure stable MD simulations, estimate phase diagrams, compute minimum energy reaction pathways, etc. Given the principium of AGNI, whereby, the force on an atom is learnt based on its environment deprives the means to predicting energy directly. Nevertheless, we now discuss some alternative strategies to estimating energy, be it in dynamic or static simulations.

**During a Molecular Dynamics Simulation.** The rate of change of the total potential energy, in an MD simulation, can be expressed as a function of the individual atomic forces and velocities by invoking the chain rule

$$\frac{dE}{dt} = \sum_{i,u} \frac{\partial E}{\partial r_i^u} \frac{\partial r_i^u}{\partial t} = - \sum_{i,u} F_i^u v_i^u \quad (5)$$

$E$  is the total potential energy of the system,  $r_i^u$  and  $v_i^u$  are the position and velocity of atom  $i$  along one of the three coordinates,  $u \subset (x, y, z)$ . In the limit that  $\partial t \rightarrow 0$ , in eq 5, an analytical expression for the rate of change in energy for infinitesimally small time difference ( $\Delta t$ ), albeit from the initial configuration, can be expressed as

$$E_t = E_{t-\Delta t} - \Delta t \left( \sum_{i,u} F_i^u v_i^u \right) \quad (6)$$

In MD simulations this translates to choosing a small time step to ensure accurate force integrations and minimize numerical noise propagation, providing a pathway to indirectly

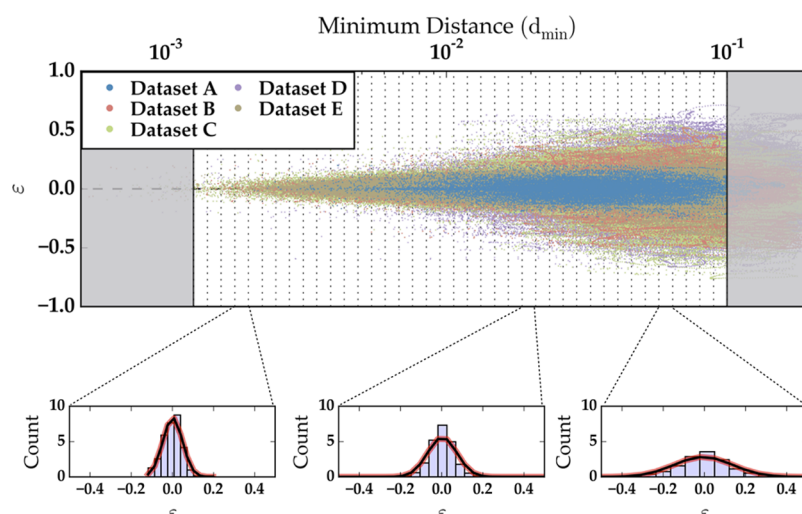
monitor the energy evolution during the course of the simulation.

To validate this scheme, constant temperature MD simulations of bulk fcc Al ( $\sim 250$  atoms) were carried out at different temperatures using the  $M_D$  AGNI force field. A time step of 0.5 fs was chosen. Using the reference forces and velocities, along with eq 6, the rate of change in energy was computed as a function of time. The results are plotted in Figure 10a. Clearly, the computed energies (from AGNI predicted forces) are conserved in time and maintain the correct ordering as a function of temperature. Note that the change in energy is reported with reference to the starting configuration. For comparison DFT energies (and scaled with respect to the starting configuration) for snapshots of atomic configurations along the dynamic trajectory are also reported. The error between the computed and ab initio predicted energies is  $\sim 4$  meV/atom, close to the order of numerical noise that one can expect in such simulations. For completeness, we also report the change in energy for MD simulations run with an EAM force field (Figure 10b). All three methods display similar trends suggesting that the dynamics undertaken with an AGNI force field does indeed concur with the thermodynamic driving forces in a system.

**During a Static Simulation.** A second, and equally simple, approach to estimating the potential energy directly from forces is by integrating them using a Taylor series approximation of the potential energy

$$E = E_0 + \sum_{i,u} \left( \frac{dE}{dr_i^u} \Big|_{r_i^u=r_o} (r_i^u - r_o) + \frac{1}{2!} \frac{d^2E}{dr_i^{u2}} \Big|_{r_i^u=r_o} (r_i^u - r_o)^2 + \dots \right) \\ \cong E_0 - \sum_{i,u} F_i^u \Delta r_i^u, \quad u \subset (x, y, z) \quad (7)$$

Here,  $E$  is once again the total potential energy, which to a first order approximation can be derived from the atomic forces.  $r_i^u$  is once again the atomic position. A discretization along the atomic positions governs the accuracy by which we can predict the energy. For these reasons eq 7 is particularly more suited for static simulations, e.g., computing reaction barriers along a reaction coordinate, when the velocities are zero. The validity of this approach was demonstrated to be consistent with the underlying potential energy surface in our previous work,



**Figure 11.** Top panel: a scatter plot of the minimum distance ( $d_{\min}$ ) vs the predicted force error ( $\epsilon$ ). The range of  $d_{\min}$  is further subdivided into small groups for statistical analysis. The gray regions were not considered for any statistical purposes, due to the lack of sufficient data (left) and high errors (right). Bottom panel: a standard normal distribution fit for each subgroup (though only shown for three such bins), used to estimate the variance in model errors.


whereby, the migration energy for a vacancy in bulk Al was within 3% of the DFT predicted value.<sup>17</sup>

Equations 6 and 7 both provide a restricted means to computing the energy, whereby, it is necessary that a pathway connecting the different configurations in phase space (either in time or along a reaction coordinate) exists, in order to accurately carry out force integration. This is a limitation of using a truly force based force field, wherein, one cannot predict energies by simply choosing two arbitrary points in the phase space. Nevertheless, these findings ascertain that eqs 6 and 7 can indeed be used to compute change in the total potential energy as a function of time or reaction coordinate, as needed by a majority of atomistic simulations.

**Uncertainty Quantification.** The final component to a successful predictive model is to be able to quantify the error ( $\epsilon$ ) in the force predictions made. If this uncertainty can be estimated, a priori, confidence estimates for the force predictions given a new atomic environment can be provided. At the same time, it allows one to understand the force field's domain of applicability, providing a pathway for their subsequent and continual improvement. Below is one such attempt to quantify these uncertainties.

To compute the force on an atom, within the learning framework (cf., eq 3), begins by calculating the distance between its fingerprint and the reference training fingerprints, resulting in a total of  $N_t$  distances. The final prediction is then a weighted sum of these distances, making them an important metric on predictive accuracy. Among the list of distances the minimum distance,  $d_{\min} = \min \{d_1, d_2, \dots, d_{N_t}\}$ , in particular provides a measure of “closeness” of the new observation compared to the reference cases and can be thought of as a descriptor in estimating  $\epsilon$ . To capture this hypothesis, the  $d_{\min}$  and  $\epsilon$  for every observation in the test data set was computed with the constructed AGNI force field,  $M_D$ . The results are summarized in the scatter plot of Figure 11. Clearly, as  $d_{\min}$  increases the variance in  $\epsilon$  increases, suggesting that for instances far away from reference training environments the interpolative performance fails. By binning the range of  $d_{\min}$ 's observed into uniform and smaller subgroups, a standard normal distribution is fit to the observed  $\epsilon$ . The histogram

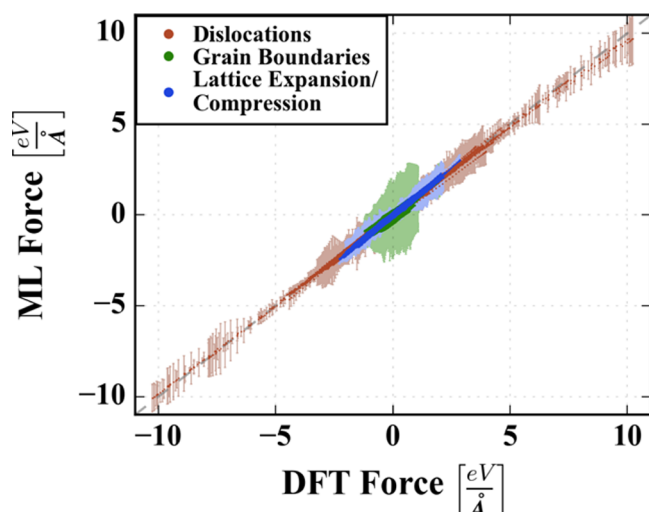
insets in Figure 11 demonstrate this for three such bins. During the binning process,  $d_{\min} < 10^{-3}$  and  $> 10^{-1}$  were ignored due data scarcity and large predictive errors, respectively (cf., gray-shaded regions in Figure 11). Collecting statistics across the range  $d_{\min}$ 's allowed us to fit a polynomial relation between standard deviation ( $s$ ) and  $d_{\min}$ , providing an analytical form to quickly estimating uncertainties. The exact functional form is  $s = 49.1d_{\min}^2 - 0.9d_{\min} + 0.05$ . This is illustrated by the red circle markers and the blue dashed line in Figure 12. Note that by



**Figure 12.** Uncertainty model, created for force field  $M_D$ , whereby  $d_{\min}$  is used as a descriptor to measure the expected variance in the prediction made. The markers show the actual behavior, while the blue dashed line indicates a polynomial fit to the uncertainty.

using  $s$  the confidence level in the uncertainties provided is at 68.2%, though one can use higher confidence levels such as  $2s$  and beyond, a call that the user needs to make depending on the need and availability of computational resources.

Now to demonstrate this uncertainty model, for each atom in the validation configurations (grain boundaries, lattice expansion/compression, and dislocation), we estimate the uncertainty in the force predictions made. In Figure 13 we replot the reference DFT and ML force predictions, along with the corresponding uncertainty in each prediction as highlighted by the error bars (color coded according to the configuration subclass marker). Immediately, those atomic environments with high uncertainties is evident and can now be flagged. These environments can be accumulated and used to retrain the force



**Figure 13.** Parity plot of force predictions for test environments (grain boundaries, lattice expansion/compression, and dislocation) of data set E, predicted by DFT and the ML model,  $M_D$ . These values are identical to those in Figure 7. For each prediction the uncertainty is highlighted, as indicated by the error bars (color coded according to the configuration subclass).

field. The preliminary steps undertaken here allow for identifying regions of poor force field performance in a systematic manner and are integral to the continual improvement in accuracy and generalizability of AGNI force fields, making them truly adaptive.

## CONCLUSIONS

A new machine learning framework to circumvent the accuracy, cost, and generalizability issues facing current atomistic models has been proposed. By directly mapping quantum mechanical derived force components to the numerical representation of an atom's local environment, accurate and computationally inexpensive force fields, herein called AGNI, were developed. In this manuscript a workflow for their systematic construction, which includes generating reference data, representing the atomic environments with a numerical fingerprint, sampling nonredundant data, learning the forces, were all demonstrated for the example of elemental Al. Further, methods to quantify uncertainties in the force predictions are proposed. This is crucial to understanding the domain of applicability of such data-driven methods, in turn paving the way for their adaptive refinement.

Nevertheless, to make such methods a mainstream tool for atomistic simulations a few challenges yet remain that need to be addressed. First, to explore diverse chemistries it is necessary to come up with AGNI force fields for multielemental systems in an equally quick and rational manner. Though the framework discussed here was for an elemental system, the recipe is directly transferable to multielemental situations. Second, as materials science or chemical systems become ever increasingly complex, the configuration space to be explored will increase exponentially. This poses a challenge for the nonlinear regression learning algorithm proposed here, and for a continued realization of machine learning force fields adopting methodologies, wherein, large quantities of data can be handled will be required.

Irrespective of these challenges, the prospect of using AGNI force fields as a tool to accelerate atomistic simulations is

indeed very promising. Access to such high fidelity force predictions at a fraction of the cost has already made significant inroads to studying materials and chemical phenomena. Our previous work demonstrated an expose of some atomistic simulations, such as geometry optimization of atomic structures with several 100s of atoms, dynamical evolution of defects over long time scales (vacancies and adatoms) to determine diffusion barriers, computing vibrational properties of materials, and estimating reaction energy barriers, all using AGNI force fields.<sup>17</sup> Here, we further extended the scope of such force fields by simulating even more complex phenomena, estimating the melting and stress–strain behavior of Al surfaces. Also, methods to reconstruct energies entirely from forces were proposed. The force field construction workflow put in place here allowed us to study these more complex materials and chemical phenomena, and such strategies are only going to become increasingly important in pushing the envelope of atomistic simulations.

## ASSOCIATED CONTENT

### Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: 10.1021/acs.jpcc.6b10908.

All quantum mechanical reference data files used in training the force field are provided in the *extended-xyz* format contained in a zip file, via the publishers Web site. The data is arranged in 5 different folders, each corresponding to a different data set (A, B, C, D or E) used in this work. Note that each folder only contains new environments to save memory. For instance, data set C should contain the configurations within the folders of A and B, as well as C. Readers can also access the configurations from our research group's data repository (<http://khazana.uconn.edu>). Also contained in the zip file are LAMMPS source code files, which can be used to run MD simulations using the new class of AGNI force fields, along with a reference force field file for Al. (ZIP)

## AUTHOR INFORMATION

### Corresponding Author

\*E-mail: [botuvenkatesh@gmail.com](mailto:botuvenkatesh@gmail.com).

### ORCID

V. Botu: 0000-0002-6244-4817

R. Batra: 0000-0002-1098-7035

### Notes

The authors declare no competing financial interest.

## ACKNOWLEDGMENTS

This work was supported financially by the Office of Naval Research (Grant No. N00014-14-1-0098). The authors would like to acknowledge helpful discussions with K. B. Lipkowitz, S. K. Yadav, G. Pilania, and A. Mannodi-Kanakkithodi. Partial computational support through an Extreme Science and Engineering Discovery Environment (XSEDE) allocation is also acknowledged.

## REFERENCES

- (1) Tadmor, E. B.; Miller, R. E. *Modeling Materials: Continuum, Atomistic and Multiscale Techniques*; Cambridge University Press: Cambridge, U.K., 2012; pp 153–300.

- (2) Hautier, G.; Jain, A.; Ong, S. P. From the computer to the laboratory: materials discovery and design using first-principles calculations. *J. Mater. Sci.* **2012**, *47*, 7317–7340.
- (3) Burke, K. Perspective on density functional theory. *J. Chem. Phys.* **2012**, *136*, 150901.
- (4) Neugebauer, J.; Hickel, T. Density functional theory in materials science. *Comp. Mol. Sci.* **2013**, *3*, 438–448.
- (5) Hill, J.-R.; Freeman, C. M.; Subramanian, L. *Use of Force Fields in Materials Modeling*; John Wiley and Sons, Inc.: New York, 2007; pp 141–216.
- (6) Torrens, I. M. *Interatomic potentials*; Academic Press Inc.: New York, 1972.
- (7) Elliott, J. A. Novel approaches to multiscale modeling in materials science. *Int. Mater. Rev.* **2011**, *56*, 207.
- (8) Bianchini, F.; Kermodé, J. R.; Vita, A. D. Modelling defects in Ni-Al with EAM and DFT calculations. *Modell. Simul. Mater. Sci. Eng.* **2016**, *24*, 045012.
- (9) Witten, I. H.; Frank, E.; Hall, M. A. *Data mining: Practical machine learning tools and techniques*; Elsevier: Amsterdam, 2011.
- (10) Hastie, T.; Tibshirani, R.; Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed.; Springer: New York, 2009.
- (11) Hofmann, T.; Scholkopf, B.; Smola, A. J. Kernel methods in machine learning. *Ann. Statist.* **2008**, *36*, 1171.
- (12) Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **2011**, *134*, 074106.
- (13) Bartók, A. P.; Payne, M. C.; Kondor, R.; Csányi, G. Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons. *Phys. Rev. Lett.* **2010**, *104*, 136403.
- (14) Lorenz, S.; Groß, A.; Scheffler, M. Representing high-dimensional potential-energy surfaces for reactions at surfaces by neural networks. *Chem. Phys. Lett.* **2004**, *395*, 210–215.
- (15) 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.
- (16) Botu, V.; Ramprasad, R. Adaptive machine learning framework to accelerate ab initio molecular dynamics. *Int. J. Quantum Chem.* **2015**, *115*, 1074–1083.
- (17) Botu, V.; Ramprasad, R. Learning scheme to predict atomic forces and accelerate materials simulations. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2015**, *92*, 094306.
- (18) Li, Z.; Kermodé, J. R.; De Vita, A. Molecular Dynamics with On-the-Fly Machine Learning of Quantum-Mechanical Forces. *Phys. Rev. Lett.* **2015**, *114*, 096405.
- (19) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133–A1138.
- (20) Hohenberg, P.; Kohn, W. Inhomogeneous Electron Gas. *Phys. Rev.* **1964**, *136*, B864–B871.
- (21) Car, R.; Parrinello, M. Unified Approach for Molecular Dynamics and Density-Functional Theory. *Phys. Rev. Lett.* **1985**, *55*, 2471–2474.
- (22) Kresse, G.; Furthmüller, J. Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1996**, *54*, 11169.
- (23) Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1999**, *59*, 1758.
- (24) Perdew, J. P.; Burke, K.; Wang, Y. Generalized gradient approximation for the exchange-correlation hole of a many-electron system. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1996**, *54*, 16533.
- (25) Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1994**, *50*, 17953.
- (26) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2013**, *87*, 184115.
- (27) Jolliffe, I. *Statistics Reference Online*; John Wiley and Sons, Inc.: New York, 2014.
- (28) Schölkopf, B.; Smola, A.; Müller, K.-R. Kernel principal component analysis. International Conference on Artificial Neural Networks, 1997; pp 583–588.
- (29) Cox, T. F.; Cox, M. A. *Modern multidimensional scaling*; Chapman and Hall/CRC: London, 2001.
- (30) Behler, J. Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.* **2011**, *13*, 17930.
- (31) Müller, K. R.; Mika, S.; Ratsch, G.; Tsuda, K.; Scholkopf, B. An introduction to kernel-based learning algorithms. *IEEE Trans. Neural Networks* **2001**, *12*, 181.
- (32) Hansen, K.; Montavon, G.; Biegler, F.; Fazil, S.; Rupp, M.; Scheffler, M.; von Lilienfeld, O. A.; Tkatchenko, A.; Müller, K. Assessment and validation of machine learning methods for predicting molecular atomization energies. *J. Chem. Theory Comput.* **2013**, *9*, 3404.
- (33) Rupp, M. Machine learning for quantum mechanics in a nutshell. *Int. J. Quantum Chem.* **2015**, *115*, 1058–1073.
- (34) Jacobsen, K. W.; Norskov, J. K.; Puska, M. J. Interatomic interactions in the effective-medium theory. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1987**, *35*, 7423–7442.
- (35) Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.* **1995**, *117*, 1–19.
- (36) Alavi, S.; Thompson, D. L. Molecular Dynamics Simulations of the Melting of Aluminum Nanoparticles. *J. Phys. Chem. A* **2006**, *110*, 1518–1523.
- (37) Neyts, E. C.; Bogaerts, A. Numerical Study of the Size-Dependent Melting Mechanisms of Nickel Nanoclusters. *J. Phys. Chem. C* **2009**, *113*, 2771–2776.
- (38) Zhang, K.; Stocks, G. M.; Zhong, J. Melting and premelting of carbon nanotubes. *Nanotechnology* **2007**, *18*, 285703.
- (39) Zhou, Y.; Karplus, M.; Ball, K. D.; Berry, R. S. The distance fluctuation criterion for melting: Comparison of square-well and Morse potential models for clusters and homopolymers. *J. Chem. Phys.* **2002**, *116*, 2323–2329.
- (40) Stoltze, P.; Norskov, J. K.; Landman, U. Disorder and Melting of Aluminum Surfaces. *Phys. Rev. Lett.* **1988**, *61*, 440–443.
- (41) Pham, H. H.; Williams, M. E.; Mahaffey, P.; Radovic, M.; Arroyave, R.; Cagin, T. Finite-temperature elasticity of fcc Al: Atomistic simulations and ultrasonic measurements. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2011**, *84*, 064101.