



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Learning scheme to predict atomic forces and accelerate materials simulations

V. Botu and R. Ramprasad

Phys. Rev. B **92**, 094306 — Published 25 September 2015

DOI: [10.1103/PhysRevB.92.094306](https://doi.org/10.1103/PhysRevB.92.094306)

A learning scheme to predict atomic forces and accelerate materials simulations

V. Botu

Department of Chemical and Biomolecular Engineering, University of Connecticut, Storrs, CT 06269

R. Ramprasad

Department of Materials Science and Engineering, University of Connecticut, Storrs, CT 06269

The behavior of an atom in a molecule, liquid or solid is governed by the force it experiences. If the dependence of this vectorial force on the atomic chemical environment can be *learned* efficiently with high-fidelity from benchmark reference results—using “big data” techniques, i.e., without resorting to actual functional forms—then this capability can be harnessed to enormously speed up *in silico* materials simulations. The present contribution provides several examples of how such a *force field* for Al can be used to go far beyond the length-scale and time-scale regimes accessible presently using quantum mechanical methods. It is argued that pathways are available to systematically and continuously improve the predictive capability of such a learned force field in an adaptive manner, and that this concept can be generalized to include multiple elements.

The dynamic behavior of an atom in a molecule, liquid or solid is directly determined by the *local* force it experiences. Nevertheless, as already pointed out by Feynman¹, forces are generally viewed as secondary computed quantities and are obtained through the agency of the total potential energy—a *global* property of the entire system. In practice, forces on atoms are obtained either as by-products during a potential energy evaluation, or from the first derivative of the potential energy with respect to the atomic positions. Direct and rapid access to atomic forces, given just the atomic configuration of a system (molecule, liquid, or solid), immediately makes it possible to perform efficient geometry optimizations and molecular dynamics (MD) simulations, provided, of course, the predicted force is formally conservative, i.e., it is consistent with an underlying potential energy surface. If the capability to predict conservative forces preserves the fidelity of high-level quantum mechanics based methods, but comes at a minuscule fraction of the cost, and if this capability can be extended systematically and progressively to potentially all configurational and chemical environments that an atom may experience, we will have a powerful and adaptive materials simulation scheme.

The present contribution lays the ground work and takes initial steps towards the above vision. A new scheme is presented that systematically *learns* in an interpolative manner to predict atomic forces in environments encountered during the dynamical evolution of materials from a set of high-level calculations performed on reference atomic configurations with modest system sizes. This concept is resonant with emerging data-driven (or “big data”^{2–6}) approaches aimed at materials discovery in general^{7,8}, as well as at accelerating materials simulations^{9–13}. Machine learning (ML) methods using neural networks^{9,10} and Gaussian processes^{11,12} have been successful in the development of interatomic potentials, wherein the potential energy surface is learned from a set of higher-level (quantum mechanics based) reference calculations.

The distinctive aspect of the present contribution,

namely, learning to predict atomic forces directly from past data has been suggested only recently^{12,13} (to accelerate *ab initio* MD simulations on-the-fly). Here, we propose the creation of a stand-alone purely data-driven force prediction recipe (devoid of any explicit functional form), that can also provide the underlying potential energy surface (through integration). Both the forces and the potential energy can be predicted with a high level of accuracy at speeds several orders of magnitude faster than the reference quantum mechanics based calculations. Moreover, this *force field* is adaptive (i.e., new conformational environments can be systematically incorporated as required), and generalizable (i.e., the scheme can be extended to any collection of elements for which reliable reference calculations can be performed). A practical scheme that exploits the rapid high-fidelity force prediction capability within a materials simulation framework is presented, and demonstrated for Al in several configurational environments and dynamical situations that go well beyond the reaches of conventional first principles simulations. Pathways to extend this concept to handle multi-elemental systems are also proposed.

Central to this development is a robust scheme to numerically and simply represent, or *fingerprint*, the atomic environments. Such a fingerprint should differentiate dissimilar configurations with adequate accuracy, and be invariant to transformations of the environment such as translation, rotation and permutation of like elements. While several such prescriptions have been proposed in the past^{12–18}, the present objective, namely, mapping the vectorial force experienced by an atom to its configurational environment, places stringent constraints on the nature of the fingerprint. We argue that the following fingerprint function, $V_i^k(\eta)$, may be used to accurately represent the k^{th} component of the force on atom i ¹²:

$$V_i^k(\eta) = \sum_{j \neq i} \frac{r_{ij}^k}{r_{ij}} \cdot e^{-\left(\frac{r_{ij}}{\eta}\right)^2} \cdot f(r_{ij}). \quad (1)$$

r_{ij} is the distance between atoms i and j , while r_{ij}^k is a scalar projection of this distance along component k . To

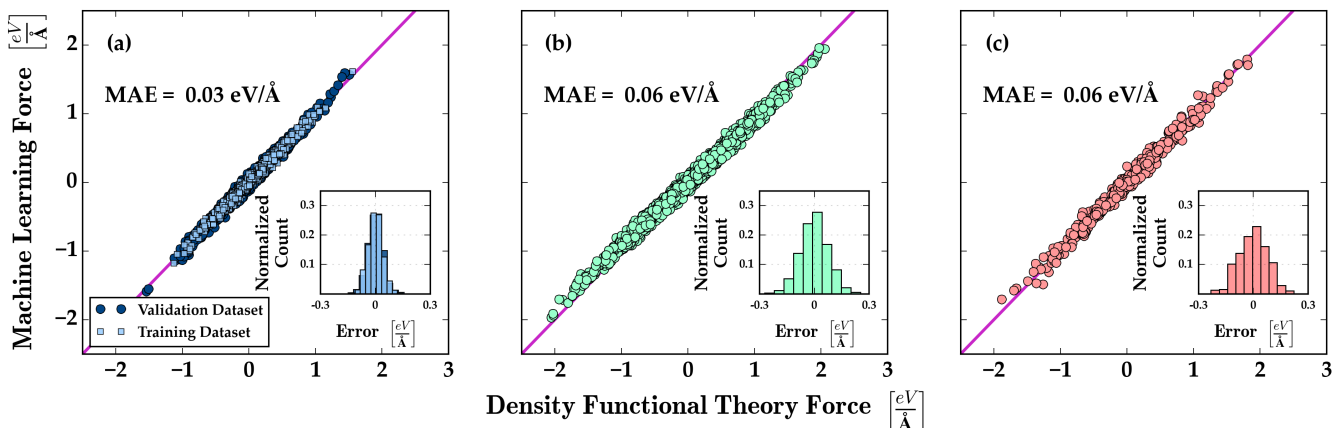


FIG. 1. Comparison of the forces predicted using the ML force field with reference DFT results, for (a) the trained model (light blue) and the validation dataset (dark blue), (b) a test unit cell containing over 800 Al atoms in the fcc phase, and (c) a test unit cell containing over 160 atoms in a hypothetical bcc phase. In (b) and (c), atoms were randomly perturbed from their equilibrium positions. Insets show the distribution of the prediction errors (defined as the difference between the predicted and reference DFT forces) leading to respectable mean absolute errors (MAE).

determine the force on an atom, we require three such components along non-parallel directions. The parameter η governs the extent of co-ordination around atom i that needs to be captured. The fingerprint is essentially a spectrum of V_i^k values corresponding to predetermined choices of η values, i.e., V_i^k is defined in an η -grid. The diminishing influence of faraway atoms is handled by a damping function, $f(r_{ij}) = 0.5 \left[\cos\left(\frac{\pi r_{ij}}{R_c}\right) + 1 \right]$. The summation in Eq. 1 runs over all neighboring atoms within an arbitrarily large cutoff distance R_c (8 Å, in the present work). By construction, the fingerprint is symmetry-adapted, and respects asymmetries of potential wells. For instance, due to the first term in the summation of Eq. 1, an atom in a centro-symmetric position will lead to a fingerprint with all zero values (to be mapped to a zero force), and an atom displaced from a centro-symmetric position will lead to fingerprints with non-zero components whose values will depend on the magnitude and direction of the off-center displacement.

The next step is to map the fingerprints to appropriate force components. Here, we have adopted the kernel ridge regression (KRR) method, capable of handling complex non-linear relationships^{12,17,18}. The KRR method works on the principle of similarity. By comparing an atom's fingerprint, $V_i^k(\eta)$, with a set of reference cases, an interpolative prediction of the k^{th} component of the force (F_i^k) can be made, and is given by

$$F_i^k = \sum_t \alpha_t \cdot \exp\left[-\frac{\|V_i^k(\eta) - V_t^k(\eta)\|^2}{2\sigma^2}\right]. \quad (2)$$

t labels each reference atomic environment, and $V_t^k(\eta)$ is its corresponding fingerprint. $\|V_i^k(\eta) - V_t^k(\eta)\|$ is the Euclidean distance between the two atomic fingerprints, though other distance metrics can be used. α_t s and σ are the weight coefficients and length scale parameter,

respectively. The optimal values for α_t s and σ are determined during the training phase, with the help of cross-validation and regularization methods^{12,17,19,20}. Further details concerning the fingerprint construction and the learning algorithm can be found elsewhere¹².

Using the above prescribed framework, a ML force field for Al has been developed using a plethora of reference atomic environments accumulated from density functional theory (DFT) based MD runs at various temperatures using the Vienna ab initio simulation package^{21–24} (other means may also be used to generate the reference data, as long as they satisfy prescribed demands on accuracy of the atomic forces). To ensure a diverse set of reference cases, Al in different geometric arrangements were considered (but each one with just a few tens of atoms per repeating unit cell), including defect-free bulk in the face centered cubic (fcc) phase, bulk fcc phase with vacancy, clean (111) surface, and the (111) surface with adatom, resulting in over 100,000 atomic environments¹². Interestingly, a random set of 1000 atomic environments drawn from the accumulated environments proved sufficient to construct an accurate interpolative force prediction model. Figure 1(a) compares the predicted forces with the DFT forces (including the error distribution in the inset) for all accumulated configurations, i.e., those used in the training phase and the remaining configurations whose results were used for validation. The mean absolute error (MAE) of the prediction model was 0.03 eV/Å, of the order of the expected chemical and numerical accuracy of the reference DFT calculations. Furthermore, this procedure to predict atomic forces is also extremely expedient; it scales linearly with system size, and can be well over 8 orders of magnitude faster than a typical DFT calculation.

An immediate (and straight-forward) application of this fast high-fidelity capability to predict atomic forces




FIG. 2. Arrhenius plots for (a) vacancy migration in bulk Al and (b) adatom diffusion on the Al (111) surface. For each temperature, the MD simulation time was extended so as to allow at least 50 hopping events (thus allowing estimation of an average hop rate, and the indicated error bar). A linear fit (solid red line) was used to determine the dynamic activation energy (E_a), and is compared with the static DFT activation energy (indicated in brackets as ‘Ref’). (c) For the vacancy migration in bulk Al, the DFT potential energy along the migration trajectory (symbols and dashed line), and the corresponding energy obtained via an integration of the ML forces along the reaction coordinate (solid line).

is geometry optimization, including the prediction of potential energy minima and saddle points. Simulations involving hundreds of thousands of atoms (i.e., cases that are beyond the reaches of present day DFT computations) can be handled, provided the chemical environments encountered during the course of such optimizations are included in the force field. In order to understand the limits of the constructed ML force field for Al within the context of such simulations, a few tests were performed. The first one involved a large unit cell containing over 800 Al atoms in the fcc phase along with Al vacancies. Atoms were randomly perturbed, and the ML force field was used to optimize this perturbed structure. The correct equilibrium geometry was recovered, as ascertained by a separate DFT calculation starting with the same perturbed system. A video of this optimization is included in the Supplemental Information²⁵. Figure 1(b) compares the predicted forces with the DFT forces for the initial perturbed geometry. Although we restrict ourselves to modest sizes in this discussion (as we are constrained by the inability of DFT to provide validation for truly large unit cells), this example demonstrates that the force field is transferrable to much larger systems, thus going significantly beyond previous efforts^{12,13}.

As a second geometry optimization example, a 160 atom unit cell of Al in a hypothetical body centered cubic (bcc) phase was considered. Once again, the atoms were perturbed randomly, followed by geometry optimization. Figure 1(c) captures the performance of the force field for the starting geometry. Given that the bcc phase was never used in the training phase during the force field creation, we would expect that forces on atoms in such an environment will be difficult to predict. Surprisingly, going by the rather small force error distribution (comparable to the 800-atom fcc example), we conclude that the

current choice of fingerprints allows us to effectively capture diverse chemical environments in a versatile manner.

Next, we consider non-zero temperature dynamical situations. For the force prescription to correctly capture dynamic processes with high-fidelity, ergodicity has to be preserved. In other words, the average behavior and time scales of elementary steps or processes should be correctly represented during a MD simulation using the force field. As a first example, we consider the diffusion of an Al vacancy in bulk Al, using a unit cell containing 32 Al sites and an Al vacancy. MD simulations were performed at 9 temperatures in the 500-900 K range for times up to 5 ns, with a timestep of 0.5 fs. By observing the dynamics of the vacancy, the average rate constant (k) for the migration process at each temperature was determined. k is given as $1/t_{hop}$, where t_{hop} is the average time taken for a vacancy to migrate to a neighboring site. To ensure that sufficient statistics are collected, k was averaged over 50 such hop events at each temperature. Figure 2(a) shows an Arrhenius plot of k versus the reciprocal temperature, whose slope yields the activation energy (E_a) for Al vacancy migration to be 0.49 eV. The corresponding DFT value for a similar, but static, migration process was determined to be 0.59 eV (c.f., Figure 2(c)). Barrier “softening” is expected under dynamical conditions, relative to the results of static calculations in which entropic effects are neglected^{26,27}.

Another elementary process we considered was the self-diffusion of an Al adatom on the Al (111) surface, using a $6 \text{ \AA} \times 6 \text{ \AA}$ surface unit cell containing a 4-layer thick Al slab. Similar to the Al vacancy example, by monitoring the dynamics of the adatom across a temperature range of 50-300 K, an E_a of 0.03 eV was predicted, as shown in Figure 2(b), whilst a static DFT calculation yielded an E_a of 0.04 eV. A video of the adatom migration MD sim-

ulation is included in the Supplemental Information²⁵.

Both dynamical diffusion scenarios considered lead to the correct Arrhenius behavior indicating that the underlying physics is properly captured in the ML force field based MD simulations. Moreover, although the force field is aimed at directly predicting atomic forces, potential energy differences for elementary steps may be obtained by integrating the forces along a suitable reaction coordinate. Figure 2(c), for instance, portrays the DFT energy profile along the Al vacancy migration pathway in bulk Al, as well as the corresponding energy determined by integrating the forces predicted by the ML force field. The close agreement between the two energies is self-evident, indicating that energies corresponding to critical parts of a trajectory may indeed be obtained from the forces through integration. More importantly, this demonstration places the force prediction scheme in a formally solid framework as the predicted ML forces are shown to be consistent with the underlying potential energy.

Lastly, we evaluate the prospect of how well thermal behavior of materials can be simulated using the force-based framework. In particular, we focus on the vibrational (or phonon) density of states (DOS), which has to be properly captured to allow for accurate calculations of thermodynamic quantities, thermal expansion, thermal conductivity, etc. Figure 3(a) shows the phonon band structure as determined using the ML force field and using DFT, and in both cases, the finite displacement method was used²⁸. Figure 3(b) shows the correspond-




FIG. 3. (a) Phonon band structure, (b) phonon density of states (DOS), and (c) Helmholtz free energy and constant volume heat capacity computed using the ML force field (solid lines) and DFT (dashed lines). The phonon band structure and DOS were computed using the finite atomic displacement method. Also included in (b) are the DOS results obtained from the Fourier transform of the velocity autocorrelation function (solid cyan hatched fill).

ing DOS, as well as the DOS computed using the Fourier transform of the velocity autocorrelation obtained from a MD simulation²⁹. This latter approach implicitly includes anharmonicity to all orders (the first method, in contrast, includes just the harmonic part). The MD simulation involved a 864 atom unit cell, and a simulation time of 5 ps at 700 K. Excellent agreement of the ML force field result with the reference DFT calculations can be seen. The deviations of the DOS computed using MD simulations relative to that obtained using the finite displacement scheme (especially at high frequencies) may be attributed to non-zero anharmonic effects. The DOS can be utilized to determine thermodynamic properties such as the Helmholtz free energy and the constant volume heat capacity. These properties, as a function of temperature, are compared with the corresponding DFT results in Figure 3(c). The ML force field and DFT results are nearly indistinguishable, indicating that even under the stringent test of small atomic perturbations encountered in these situations (as opposed to the larger length scale vacancy or adatom hops discussed earlier), the fidelity of the force prediction is preserved.

A natural question that arises at this point is how this force field paradigm may be extended to include multiple elements. In a multi-elemental system, the fingerprint of an atom of a given element type may be constructed to have as many parts as the number of elements in the system. Each part would represent the arrangement of atoms of a particular elemental type around the reference atom. While this scheme requires further optimization, preliminary work shows significant promise. For two binary systems, α -Al₂O₃ and monoclinic HfO₂, the force prediction based on the concatenated multi-component fingerprint prescription rivals that for the elemental Al in quality. A parity plot comparing the predicted force with the corresponding reference DFT result for each element type is shown in the Supplemental Information²⁵. Given such accuracies, extension of the proposed concept to multielemental systems appears feasible.

The discussion thus far has provided an exposé of materials simulation examples that can benefit enormously through a capability to directly and rapidly predict atomic forces with demonstrable verisimilitude. This capability learns from past reference quantum mechanical calculations of modest system sizes, but can access length-scales and time-scales significantly beyond the reaches of the reference calculations (while preserving accuracy). Examples of phenomena that can potentially be studied include transport (thermal and mass), phase transformations and chemical reactions, mechanical behavior, materials degradation and failure, etc., all within the framework of reality-mimicking non-zero temperature dynamical simulations. Widespread use of the proposed class of learning-based *force* fields will require attending to a few critical matters. These include: (i) creation of an initial compact training set of reference atomic environments appropriate for a particular materials application; and, (ii) development of a capability to

recognize a truly new atomic environment when such is encountered during the course of a simulation. The latter aspect is critical to evaluating when the force field is expected to fail, and, as importantly, to supplement the initial training set so as to make the force prediction scheme adapt, evolve and continuously improve with time. Nevertheless, these hurdles have been encountered, and addressed, in the past in many “big data” situations^{2–6}. Hence, there is reason for (cautious) optimism in the present context of high-fidelity, adaptive and generalizable atomic force fields.

ACKNOWLEDGMENTS

This work was supported financially by a grant from the Office of Naval Research (N00014-14-1-0098). The

authors would like to acknowledge helpful discussions with K. B. Lipkowitz, G. Pilania, T. D. Huan, A. Mannodi-Kanakkithodi and A. Dongare. Partial computational support through a Extreme Science and Engineering Discovery Environment (XSEDE) allocation is also acknowledged. All calculations were performed in the pythonic environment, with the atomic simulation environment, pwtools and mlpy modules^{30–32}.

-
- ¹ R. P. Feynman, *Phys. Rev.* **56**, 340 (1939).
- ² J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, and L. Brilliant, *Nature* **457**, 1012 (2009).
- ³ H. Choi and H. Varian, *Econ. Rec.* **88**, 2 (2012).
- ⁴ S. E. Hampton, C. A. Strasser, J. J. Tewksbury, W. K. Gram, A. E. Budden, A. L. Batcheller, C. S. Duke, and J. H. Porter, *Front. Ecol. Environ.* **11**, 156 (2013).
- ⁵ M. M. Gobble, *Res. Technol. Manage.* **56**, 64 (2013).
- ⁶ P. T. Metaxas and E. Mustafaraj, *Science* **338**, 472 (2012).
- ⁷ T. Mueller, A. G. Kusne, and R. Ramprasad, in *Reviews in Computational Chemistry*, edited by A. L. Parrill and K. B. Lipkowitz (Wiley, 2015).
- ⁸ L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler, *Phys. Rev. Lett.* **114**, 105503 (2015).
- ⁹ J. Behler, *Phys. Chem. Chem. Phys.* **13**, 17930 (2011).
- ¹⁰ S. Lorenz, A. Groß, and M. Scheffler, *Chem. Phys. Lett.* **395**, 210 (2004).
- ¹¹ A. P. Bartok, M. C. Payne, R. Kondor, and G. Csanyi, *Phys. Rev. Lett.* **104**, 136403 (2010).
- ¹² V. Botu and R. Ramprasad, *Int. J. Quantum Chem.* **115**, 1074 (2015).
- ¹³ Z. Li, J. R. Kermode, and A. De Vita, *Phys. Rev. Lett.* **114**, 096405 (2015).
- ¹⁴ J. Behler, *J. Chem. Phys.* **134**, 074106 (2011).
- ¹⁵ L. Yang, S. Dacek, and G. Ceder, *Phys. Rev. B* **90**, 054102 (2014).
- ¹⁶ A. P. Bartok, R. Kondor, and G. Csanyi, *Phys. Rev. B* **87**, 184115 (2013).
- ¹⁷ M. Rupp, A. Tkatchenko, K. R. Muller, and O. A. von Lilienfeld, *Phys. Rev. Lett.* **108**, 058301 (2012).
- ¹⁸ G. Pilania, C. Wang, X. Jiang, S. Rajasekaran, and R. Ramprasad, *Sci. Rep.* **3**, 2810 (2013).
- ¹⁹ K. Hansen, G. Montavon, F. Biegler, S. Fazil, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K. Muller, *J. Chem. Theory Comput.* **9**, 3404 (2013).
- ²⁰ T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. (Springer, New York, 2009).
- ²¹ G. Kresse and J. Furthmuller, *Phys. Rev. B* **54**, 11169 (1996).
- ²² G. Kresse and D. Joubert, *Phys. Rev. B* **59**, 1758 (1999).
- ²³ J. P. Perdew, K. Burke, and Y. Wang, *Phys. Rev. B* **54**, 16533 (1996).
- ²⁴ P. E. Blöchl, *Phys. Rev. B* **50**, 17953 (1994).
- ²⁵ See Supplemental Material at [URL will be inserted by publisher] for (i) a video of geometry optimization of defects in bulk fcc Al, (ii) a video of a molecular dynamics simulation of adatom migration on the Al (111) surface, and (iii) a discussion on using the direction-resolved atomic fingerprint for multi-elemental systems, including some preliminary force predictions.
- ²⁶ V. K. La Mer, *J. Chem. Phys.* **1**, 289 (1933).
- ²⁷ S. W. Benson, *Thermochemical kinetics: methods for the estimation of thermochemical data and rate parameters*, 2nd ed. (Wiley, New York, 1976).
- ²⁸ D. Alfe, *Comp. Phys. Comm.* **180**, 2622 (2009).
- ²⁹ J. M. Dickey and A. Paskin, *Phys. Rev.* **188**, 1407 (1969).
- ³⁰ S. R. Bahn and K. W. Jacobsen, *Comput. Sci. Eng.* **4**, 56 (2002).
- ³¹ S. Schmerler, “pwtools,” <http://elcorto.bitbucket.org/pwtools/index.html> (2015).
- ³² D. Albanese, R. Visintainer, S. Merler, S. Riccadonna, G. Jurman, and C. Furlanello, “mlpy: Machine learning python,” (2012), arXiv:1202.6548.