On-the-Fly Bayesian Active Learning of Interpretable Force-Fields for Atomistic Rare Events

Jonathan Vandermause,[1,2] Steven B. Torrisi,[1] Simon Batzner,[2,3] Alexie M. Kolpak,[4] and Boris Kozinsky[2,5]

[1]*Department of Physics, Harvard University, Cambridge, MA 02138, USA*
[2]*John A. Paulson School of Engineering and Applied Sciences,*
Harvard University, Cambridge, MA 02138, USA
[3]*Center for Computational Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[4]*Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[5]*Bosch Research, Cambridge, MA 02139, USA*
(Dated: April 4, 2019)

Training machine learning based interatomic potentials often requires thousands of first principles calculations, severely limiting their practical application. We present an adaptive Bayesian inference method for automating and accelerating the on-the-fly construction of accurate interatomic force fields from a single molecular dynamics simulation. Within an online active learning algorithm, the internal uncertainty of a Gaussian process regression model is used to decide whether to accept the model prediction or to perform a first principles calculation. This approach provides several orders of magnitude acceleration of molecular dynamics simulations dominated by rare events, using minimal amounts of first-principles data to rapidly and automatically train highly accurate multi-body force fields that are competitive with state-of-the-art classical and machine learned potentials.

Recent machine learning (ML) approaches to modeling the Born-Oppenheimer potential energy surface (PES) have been shown to approach first principles accuracy for a number of molecular and solid-state systems [1–4]. These methods provide a promising path toward fast, accurate, and large-scale materials simulations with the accuracy of density functional theory (DFT) and the computational efficiency of classical molecular dynamics (MD). However, most ML potentials return point estimates of the quantities of interest (typically energies, forces, and stresses) rather than a predictive distribution that reflects model uncertainty. Without uncertainty estimates, a laborious fitting procedure is required, which usually involves selecting thousands of reference structures from a database of first principles calculations. At test time, lack of predictive uncertainty makes it difficult to determine when the fitted model is trustworthy, leading to unreliable results and lack of guidance on how to update the model in the presence of new data.

Here, we show that active learning based on Gaussian process (GP) regression can accelerate the training of high-quality force fields by making use of accurate internal estimates of model error. By combining DFT with GP regression in a single molecular dynamics simulation, an accurate multi-phase force field for bulk aluminum is obtained with fewer than 100 DFT calls. Moreover, we demonstrate that the model can be flexibly and automatically updated when the system deviates from previous training data. Such a reduction in the computational cost of training and updating potentials promises to extend ML modeling to a much wider class of materials than has been possible to date. The method is shown to successfully model rapid crystal melts and rare diffusive events, and so we call our method FLARE: Fast Learning of Atomistic Rare Events, and make the software freely available online [5].

The key contribution of this work is the development of a fully interpretable low-dimensional and non-parametric regression model of the PES that provides trustworthy estimates of model uncertainty. Typical ML schemes for modeling the PES involve regression over a high-dimensional descriptor space chosen either on physical grounds [6, 7] or learned directly from *ab initio* data [4, 8]. These approaches require building highly flexible models with many physically non-interpretable parameters, complicating the task of inferring a posterior distribution over possible models. We instead bypass the need for a high dimensional descriptor by imposing a physical prior that constrains the model to two- and three-body interactions. Because the low-dimensional descriptor space of our models can be sampled with a small amount of training data, our method avoids sparsification, a procedure that is used in Gaussian approximation potentials (GAP) to make inference tractable with many-body descriptors like SOAP [7, 9, 10], but that is also known to degrade the quality of GP variance estimates [11]. The learning task is simplified as a result, making it possible to automatically tune the model's hyperparameters in a data-driven fashion and derive trustworthy estimates of model uncertainty. This opens the door to a practical uncertainty-driven scheme for selecting training points "on-the-fly" [12], allowing an accurate potential to be learned with a minimal number of relatively expensive first principles calculations.

To reason about model uncertainty, we use *ab initio* force data to construct GP models, an established Bayesian approach to describing prior and posterior distributions over unknown functions [11]. The training database of the GP is populated with individual atomic environments by expressing the total energy of the sys-
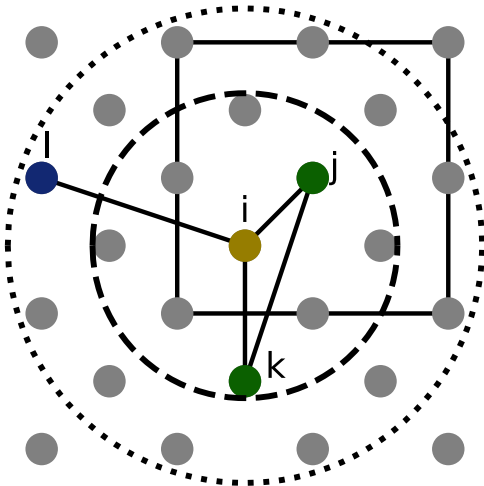
FIG. 1. The two- and three-body cutoffs used in this work. The atomic environment of the central atom (gold) consists of all position vectors of atoms within the two- and three-body cutoff spheres (dashed and dotted, respectively), including images of atoms in the primary periodic cell (solid square).

tem as a sum over two- and three-body terms,

$$E = \sum_{il} \varepsilon_{il} + \sum_{ijk} \varepsilon_{ijk}, \qquad (1)$$

where the sums range over all unique pairs and triplets of atoms that contain at least one atom from the primary periodic cell. In practice, the sums are truncated by considering local atom-centered environments $\rho = \{\vec{r}_1, ..., \vec{r}_N\}$, defined as the set of position vectors of all atoms within a chosen cutoff distance r_{cut} from the central atom (see Fig. 1). As in [13], the covariance between pair and triplet energies is defined by a kernel that directly compares interatomic distances in a rotationally invariant fashion, so that the local energy kernel between two atomic environments ρ_1, ρ_2 takes the form

$$k(\rho_1, \rho_2) = \sum_{\substack{i \in \rho_1 \\ j \in \rho_2}} k_2(d_i, d_j) + \sum_{\substack{i_1, i_2 \in \rho_1 \\ j_1, j_2 \in \rho_2}} k_3(\mathbf{d}_{i_1, i_2}, \mathbf{d}_{j_1, j_2}), \quad (2)$$

where k_2 and k_3 are the two- and three-body kernel contributions, d_i is the distance from atom i to the central atom of atomic environment ρ_1 and $\mathbf{d}_{i_1, i_2} = (r_{i_1}, r_{i_2}, r_{i_1, i_2})$ is a vector of interatomic distances of atoms i_1, i_2, and the central atom of environment ρ_1. The resulting force kernel is obtained by differentiating with respect to the Cartesian coordinates α, β of the central atoms of ρ_1 and ρ_2,

$$k_{\alpha, \beta}(\rho_1, \rho_2) = \frac{\partial^2 k(\rho_1, \rho_2)}{\partial \vec{r}_{1\alpha} \partial \vec{r}_{2\beta}}, \qquad (3)$$

giving an exactly rotationally covariant and energy conserving model [3, 13, 14]. While in this work we restrict

our attention to two- and three-body models, which are sufficiently expressive to describe both the solid and liquid phases of aluminum with high accuracy, it is straightforward to extend the kernel to capture interactions of arbitrary order, but at a computational cost that grows combinatorially with the order of the kernel.

For the pair and triplet kernels $k_{(2,3)}$, we choose the squared exponential kernel multiplied by a smooth quadratic cutoff function that ensures the potential is continuous as atoms enter and exit the cutoff sphere,

$$k_2 = \sigma_{s,2}^2 \exp\left(-\frac{(d_i - d_j)^2}{2\ell_{(2,3)}^2}\right) f_{\mathrm{cut}}(d_i, d_j),$$

$$k_3 = \sigma_{s,3}^2 \exp\left(-\frac{||\mathbf{d}_{i_1, i_2} - \mathbf{d}_{j_1, j_2}||^2}{2\ell_{(2,3)}^2}\right) f_{\mathrm{cut}}(\mathbf{d}_{i_1, i_2}, \mathbf{d}_{j_1, j_2}),$$

$$(4)$$

where $\sigma_{s,(2,3)}$ is the signal variance related to the maximum uncertainty of points far from the training set, $\ell_{(2,3)}$ is a hyperparameter that sets the length scale of the two- and three-body kernel contributions, and $||.||$ denotes the vector 2-norm. The three-body kernel is summed over all permutations of the elements of the second descriptor vector $\vec{d}_2$ in order to guarantee permutational invariance of the force model. The force $\vec{f}_i$ on each atom i and its corresponding predictive variance $\mathcal{V}[\vec{f}_i]$ are computed using the standard GP relations [11],

$$f_{i\alpha} = \bar{k}_{i\alpha}^T \left(K + \sigma_n^2 I\right)^{-1} \bar{y}$$
$$\mathcal{V}[f_{i\alpha}] = k_{\alpha, \alpha}(\rho_i, \rho_i) - \bar{k}_{i\alpha}^T \left(K + \sigma_n^2 I\right)^{-1} \bar{k}_{i\alpha}, \qquad (5)$$

where $\bar{k}_{i\alpha}$ is the vector of force kernels between ρ_i and the atomic environments in the training set, i.e. $\bar{k}_{i\alpha, j\beta} = k_{\alpha, \beta}(\rho_i, \rho_j)$, K is covariance matrix $K_{m\alpha, n\beta} = k_{\alpha, \beta}(\rho_m, \rho_n)$ of the training points, $\bar{y}$ is the vector of forces acting on the atoms in the training set, and σ_n is a hyperparameter that characterizes observation noise.

To justify an on-the-fly learning algorithm, we first characterize the uncertainty and noise estimates of the GP models and compare them against test errors on out-of-sample structures. In all models in this work, the hyperparameters $\sigma_2, \sigma_3, \ell_2, \ell_3$, and σ_n are optimized with the BFGS algorithm by maximizing the likelihood of the training data. Computation of the likelihood in GP regression involves inverting the covariance matrix K and is efficient if the model is trained on fewer than ~ 1000 points. It is worth noting that this data-driven optimization approach stands in contrast to other GP models of the PES, in which hyperparameters are chosen heuristically [2]. Remarkably, the optimized noise parameter σ_n and the predictive variance $\mathcal{V}$ are found to provide a sensitive probe of model error. We test the relationship between internal GP error and true error by performing a set of plane-wave DFT calculations on a 32-atom supercell of FCC aluminum with atomic positions randomly
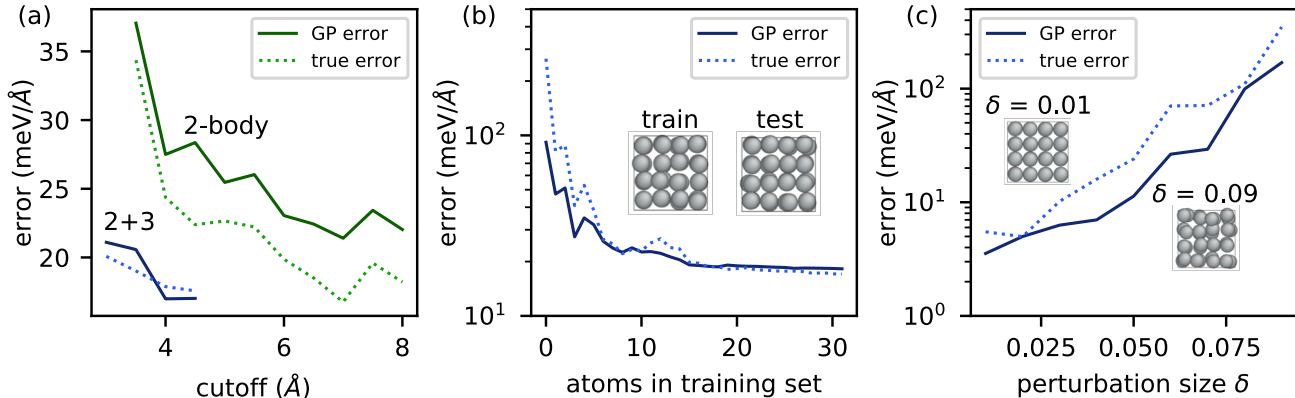
FIG. 2. Correlation of optimized Gaussian process noise σ_n and predictive variance $\mathcal{V}$ with the root mean square error on independent test structures. (a) The optimized noise parameter σ_n (solid) and root mean squared error (dotted) as a function of the cutoff radius r_{cut} of the atomic environment. (b) Combined model error $\sqrt{\sigma_n^2 + \mathcal{V}_{\text{mean}}}$ as a function of the number of training atoms. (c) Mean predictive variance $\mathcal{V}_{\text{mean}}$ (solid) and RMSE (dotted) on test structures with atomic coordinates perturbed from $\delta = 1\%$ to 9% of the lattice parameter.

perturbed from their equilibrium sites. In Fig. 2(a), all atomic coordinates are randomly perturbed by up to 5% of the lattice parameter, which was set to the experimental value of 4.046 Å. Two- and two-plus-three body GP models are trained on all forces in a single structure and tested on an independently generated structure, with the cutoff radius swept from 3.5 to 8 Å. For the two-body models, only the first term of Eq. (2) was kept, and the three body contribution was discarded. For the two-plus-three body models (lower left), the 2-body cutoff was held fixed at 6 Å and the 3-body cutoff was swept from 3 to 4.5 Å. The optimized noise parameter σ_n plotted in Fig. 2(a) closely tracks the root mean squared error (RMSE) on the test structure for the range of examined cutoff values. This provides a principled tool for selecting the cutoff radius of the GP model, showing that the expected error of the model at a given cutoff can be directly estimated from the noise hyperparameter σ_n.

When the GP model is trained on insufficient data, we find that the predictive variance $\mathcal{V}$ rises above the baseline noise level σ_n of the model, indicating that the model requires additional training data to make accurate force estimates. The utility of the predictive variance is illustrated in Fig. 2(b). Using the same training and test structures as Fig. 2(a), a GP model is constructed by adding forces on specific atoms to the training set and evaluating the RMSE and GP error after each atomic environment is added. The average GP error $\sqrt{\sigma_n^2 + \mathcal{V}_{\text{mean}}}$ is found to closely track the RMSE, where $\mathcal{V}_{\text{mean}}$ is the mean predictive variance over all atoms in the test set. We also demonstrate in Fig. 2(c) that the GP variance provides an indicator of model error when the model is forced to extrapolate on structures far from the training

set. To show this, a model was trained on a single structure with atomic coordinates perturbed by $\delta = 5\%$ of the lattice parameter and tested on structures generated with values of δ ranging from 1 to 9%. The mean variance $\mathcal{V}_{\text{mean}}$ is seen to correlate with the true error across all values of δ.

The reliability of the internal GP error estimate is the unique feature of our approach that enables us to implement FLARE, a fully adaptive active learning molecular dynamics method, in which DFT is called only when the internal error of the GP model rises above an adaptive threshold based on the optimized noise parameter σ_n. The algorithm takes an arbitrary structure as input and begins with a call to DFT, which is used to initialize a GP model. MD steps are proposed by the current GP model, with calls to DFT made whenever the error of a force prediction rises above the current noise parameter σ_n of the model, in which case the training set is augmented with the highest uncertainty atomic environment. All hyperparameters, including the noise parameter σ_n, are optimized whenever an atomic environment and its force components are added to the training set, allowing the error threshold to adapt to novel environments encountered during the simulation.

The method is implemented by coupling the Quantum ESPRESSO DFT code [15] to MD and GP code using the FLARE package [5]. We demonstrate our method by applying it to a 32-atom bulk aluminum system (Fig. 3). The simulation begins in the FCC phase at low temperature. As shown in Fig. 3(a), DFT is called often at the beginning of the simulation as the GP model learns a force field suitable for FCC aluminum. After about 30 time steps, the model needs far fewer training points,
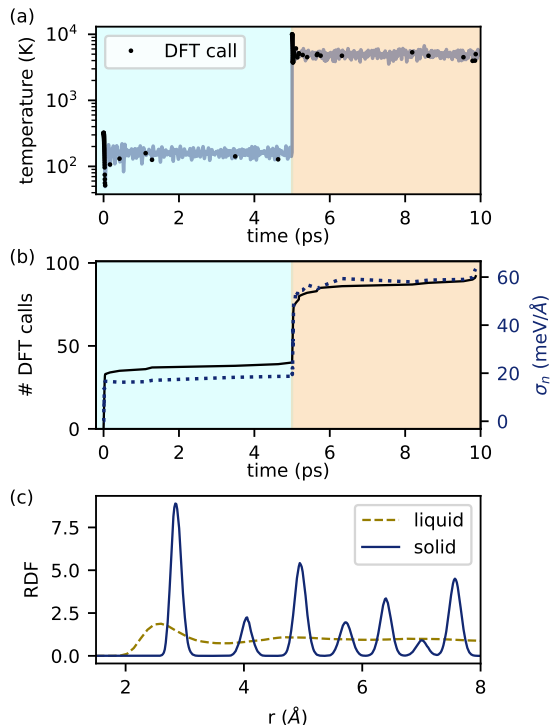
Test Set	EAM [16]	AGNI [17]	FLARE
FCC Solid	46.1	41.2	**32.9**
Liquid	157.0	128.0	**90.2**

TABLE I. Comparison of the FLARE model of Fig. 3 against other recent aluminum potentials. Test structures were drawn from *ab initio* molecular dynamics trajectories of a 2x2x2 supercell of bulk aluminum in the solid and liquid phases. Errors are reported in meV/Å, with lowest errors highlighted in bold.

phase is higher, and as a result the GP model requires a roughly similar number of DFT calls for both the solid and liquid phases. As shown in Fig. 3(b), fewer than 100 calls are needed in total during the 10 ps of dynamics, with the majority of DFT calls made at the beginning of the simulation and immediately after melting.

The performance of the obtained potential is validated by testing the model on two independent 10 ps *ab initio* molecular dynamics (AIMD) simulations of the solid and liquid phases of aluminum. 100 structures were sampled from the AIMD trajectories with 0.1 ps spacing between structures. Force predictions on all test structures were obtained with the GP potential of Fig. 3 and compared against the corresponding DFT values, with the mean absolute error in meV/Å recorded in Table I. For reference, the models are compared against a state-of-the-art aluminum EAM potential [16] and a recent aluminum ML potential [17]. Each potential is tested on the same structures, with the FLARE potential reaching the lowest force errors for both trajectories. This is partly due to the fact that the FLARE method optimizes the force field specifically for the system of interest.

Finally, we demonstrate that FLARE can be used to analyze and dramatically accelerate simulations of rare-event dynamics over timescales spanning hundreds of picoseconds by studying vacancy diffusion in a 32-atom bulk aluminum system during a 1 ns simulation. The GP model was constructed with a two-body kernel with cutoff $r_{cut} = 5.4$ Å. The system is initialized by removing one atom from an equilibrium FCC structure and setting the instantaneous initial temperature to 1500 K, giving a mean temperature of ≈ 734 K across the simulation. As shown in Fig. 4(a), most DFT calls are made early on in the simulation. After the first ~ 400 ps, no additional DFT calls are required, and the model is shown to predict vacancy hopping every few hundred picoseconds. To check the accuracy of the underlying energy model of the GP, DFT energies were computed along the high symmetry transition path shown in Fig. 4(b). The GP force predictions along the transition path were integrated to give an estimate of the energy barrier, showing close agreement to the *ab initio* DFT values. The entire FLARE run, including DFT calculations, GP training, force evaluations and MD updates, were performed on a 32-core machine in 68.8 hours of wall time. Individual

FIG. 3. Active learning of a multi-phase aluminum force field. (a) Instantaneous temperature during a 10 ps on-the-fly MD trajectory generated with FLARE. The simulation begins in the FCC phase at low temperature and is melted at $t = 5$ ps. When the predicted error on a force component rises above the current noise parameter σ_n of the model, DFT is called (black dots). (b) The number of DFT calls (solid) and optimized noise parameter (dotted) throughout the simulation. A sharp increase is observed when the crystal is melted, illustrating the model's ability to actively learn the liquid phase. (c) During the first 5 ps, the radial distribution is consistent with a low temperature solid (solid line). In the final half of the simulation, the system exhibits an RDF characteristic of the liquid phase (dashed).

requiring fewer than 50 training atoms in the first 5 ps of the simulation. To test the model's ability to actively learn and adapt to changing conditions, the crystal is melted at time $t = 5$ ps by rescaling the velocities of the atoms to give the system an instantaneous temperature of 10^4 K (well above the experimental melting point of aluminum due to the strong finite size effects of the 2x2x2 supercell). As shown in Fig. 3(b), the GP model requires a large number of DFT calls immediately after the crystal melts, as the atomic environments in the liquid phase of aluminum are significantly different from the previous training data. The noise parameter σ_n of the model sharply increases as the system enters the liquid phase, reflecting the fact that it is more difficult to model, involving more diverse atomic environments and significantly larger force fluctuations. Because the error threshold is set equal to σ_n, the threshold in the liquid
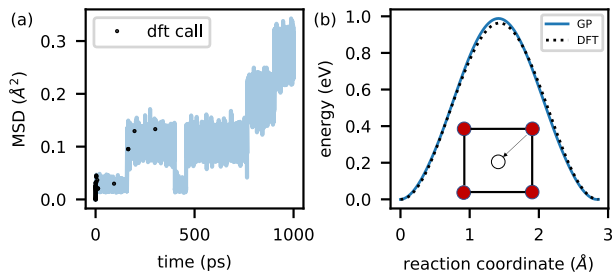
FIG. 4. Active learning of vacancy diffusion in bulk aluminum. Left: Mean squared displacement during a FLARE run of duration 1 ns. The majority of DFT calls occur at the beginning of the run, with no additional calls required after the first 400 ps. Right: the energy model of the resulting FLARE potential is tested on a high symmetry transition path, in close agreement with the *ab initio* barrier.

DFT calls required over a minute of wall time on average, making FLARE more than 300 times faster than an equivalent AIMD run.

In summary, we have presented a method for rapidly training Gaussian process models that provide highly accurate force estimates and reliable internal estimates of model uncertainty. The model's noise hyperparameter and predictive variance are shown to correlate well with the true out-of-sample error, providing a principled basis for active learning of a force field model during molecular dynamics. The FLARE non-parameteric interatomic potential model described here requires a much smaller number of atomic environments to converge the model than other state-of-the-art machine learning approaches, and is therefore well-suited to settings where large databases of *ab initio* data are too expensive to compute. Our models have a simple, accurate, and physically interpretable underlying energy model, which can be used to map the potential to a faster regression model approaching the efficiency of a classical force field [13]. This provides a path toward potentials with the accuracy of DFT at several orders of magnitude lower computational cost, which we expect to considerably expand the range of material systems that can be accurately studied with atomistic simulation. Particularly promising is the application of the FLARE engine to dynamical systems dominated by rare diffusion or reaction events, that are very difficult to treat with existing *ab initio*, classical force field or machine learning models.

[1] J. Behler, Physical Chemistry Chemical Physics **13**, 17930 (2011).
[2] V. L. Deringer and G. Csányi, Physical Review B **95**, 094203 (2017).
[3] S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, Science advances **3**, e1603015 (2017).
[4] K. Schütt, P.-J. Kindermans, H. E. S. Felix, S. Chmiela, A. Tkatchenko, and K.-R. Müller, in *Advances in Neural Information Processing Systems* (2017) pp. 991–1001.
[5] https://github.com/mir-group/flare.
[6] J. Behler, The Journal of chemical physics **134**, 074106 (2011).
[7] A. P. Bartók, R. Kondor, and G. Csányi, Physical Review B **87**, 184115 (2013).
[8] L. Zhang, J. Han, H. Wang, W. Saidi, R. Car, and E. Weinan, in *Advances in Neural Information Processing Systems* (2018) pp. 4441–4451.
[9] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, Physical review letters **104**, 136403 (2010).
[10] A. P. Bartók and G. Csányi, International Journal of Quantum Chemistry **115**, 1051 (2015).
[11] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, Vol. 2 (MIT Press Cambridge, MA, 2006).
[12] Z. Li, J. R. Kermode, and A. De Vita, Physical review letters **114**, 096405 (2015).
[13] A. Glielmo, C. Zeni, and A. De Vita, Physical Review B **97**, 184307 (2018).
[14] A. Glielmo, P. Sollich, and A. De Vita, Physical Review B **95**, 214302 (2017).
[15] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, *et al.*, Journal of physics: Condensed matter **21**, 395502 (2009).
[16] H. Sheng, M. Kramer, A. Cadien, T. Fujita, and M. Chen, Physical Review B **83**, 134118 (2011).
[17] V. Botu, R. Batra, J. Chapman, and R. Ramprasad, The Journal of Physical Chemistry C **121**, 511 (2016).