

Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons

Albert P. Bartók and Mike C. Payne

Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge, CB3 0HE, United Kingdom

Risi Kondor

Center for the Mathematics of Information, California Institute of Technology, MC 305-16, Pasadena, California 91125, USA

Gábor Csányi

Engineering Laboratory, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom

(Received 1 October 2009; published 1 April 2010)

We introduce a class of interatomic potential models that can be automatically generated from data consisting of the energies and forces experienced by atoms, as derived from quantum mechanical calculations. The models do not have a fixed functional form and hence are capable of modeling complex potential energy landscapes. They are systematically improvable with more data. We apply the method to bulk crystals, and test it by calculating properties at high temperatures. Using the interatomic potential to generate the long molecular dynamics trajectories required for such calculations saves orders of magnitude in computational cost.

DOI: 10.1103/PhysRevLett.104.136403

PACS numbers: 71.15.Nc, 31.50.-x, 34.20.Cf, 65.40.De

Atomic scale modeling of materials is now routinely and widely applied, and encompasses a range of techniques from exact quantum chemical methods [1] through density functional theory (DFT) [2] and semiempirical quantum mechanics [3] to analytic interatomic potentials [4]. The associated trade-offs in accuracy and computational cost are well known. Arguably, there is a gap between models that treat electrons explicitly and those that do not. Models in the former class are in practice limited to handling a few thousand atoms, while the simple analytic interatomic potentials are limited in accuracy, regardless of how they are parametrized. The panels in the top row of Fig. 1 illustrate the typical performance of analytic potentials in bulk semiconductors. Perhaps surprisingly, potentials that are generally regarded as adequate for describing these bulk phases show significant deviation from the quantum mechanical potential energy surface. This in turn gives rise to significant errors in predicting properties such as elastic constants and phonon spectra.

In this Letter we are concerned with the problem of modeling the Born-Oppenheimer potential energy surface (PES) of a set of atoms, but without recourse to simulating the electrons explicitly. We mostly restrict our attention to modeling the bulk phases of carbon, silicon, germanium, iron, and gallium nitride, using a unified framework. Even such single-phase potentials could be useful for calculating physical properties, e.g., the thermal expansion coefficient, the phonon contribution to the thermal conductivity, the temperature dependence of the phonon modes, or as part of hybrid schemes [5].

The first key insight is that this is actually practicable: the reason that interatomic potentials are at all useful is that the PES is a relatively smooth function of the nuclear coordinates. Improving potential modeling is difficult *not*

because the PES is rough, but because it does not easily decompose into simple closed functional forms. Secondly, away from isolated quantum critical points, the behavior of atoms is localized in the sense that if the total energy of a system is written as a sum of atomic energies,

$$E = \sum_i^{\text{atoms}} \varepsilon(\{\mathbf{r}_{ij}\}), \quad (1)$$

where $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ is the relative position between atoms i and j , then good approximations of E can be obtained by

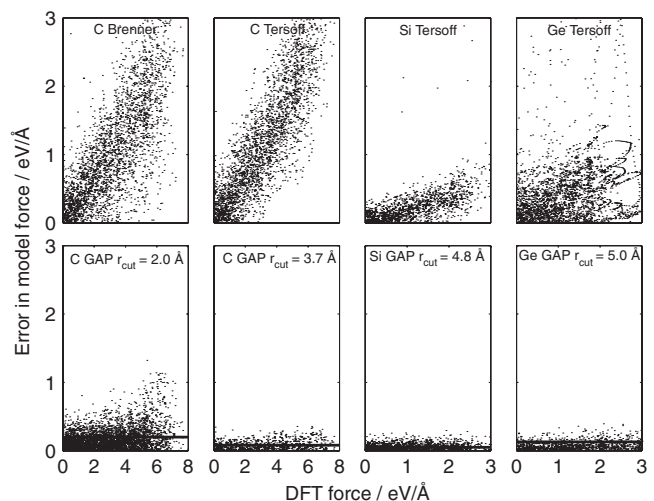


FIG. 1. Deviation of atomic forces between DFT and various models: the Brenner [18] and Tersoff [19] potentials and different GAP models for different semiconductors. In the bottom row the horizontal lines correspond to the smallest standard deviation of the error theoretically attainable given the range of the potential (see text). The configurations are taken from molecular dynamics runs at 1000 K.

restricting the set of atoms over which the index j runs to some fixed neighborhood of atom i , i.e., $|\mathbf{r}_{ij}| < r_{\text{cut}}$. In fact, we take Eq. (1) with this restriction as the defining feature of an interatomic potential. Note that in general it is desirable to separate out Coulomb and dispersion terms from the atomic energy function, $\varepsilon(\{\mathbf{r}_{ij}\})$, because the covalent part that remains can then be localized much better for the same overall accuracy. The strict localization of ε enables the independent computation of atomic energies. However, it also puts a limit on the accuracy with which the PES can be approximated. Consider an atom whose environment inside r_{cut} is fixed. The true quantum mechanical force on this atom will show a variation, depending on its environment outside the cutoff. An estimate of this variance is shown on Fig. 1 by the horizontal lines: no interatomic potential with the given cutoff can have a lower typical force error.

To date, two works have attempted to model the PES in its full generality. In the first [6], small molecules were modeled by expanding the total energy in polynomials of all the atomic coordinates, without restricting the range of the atomic energy function. While this gave extremely accurate results, it cannot scale to more than a few atoms. More recently, a neural network was used to model the atomic energy [7]. Our philosophy and aims are similar to the latter work: we compute $\varepsilon(\{\mathbf{r}_{ij}\})$ by interpolating a set of stored reference quantum mechanical results using a carefully constructed measure of similarity between atomic neighborhoods. We strive for computational efficiency in our use of expensive *ab initio* data by using both the total energy and the atomic forces to obtain the best possible estimate for ε given our assumptions about its smoothness. Furthermore, our scheme makes the generation of potential models automatic, with almost no need for human intervention in going from quantum mechanical data to the final interatomic potential model. In the following, we present an overview of our formalism. Detailed derivations are given in the supplementary information [8].

The atomic energy function is invariant under translation, rotation, and the permutation of atoms. One of the key ideas in the present work is to represent atomic neighborhoods in a transformed system of coordinates that accounts for these symmetries. Ideally, this mapping should be one-to-one: mapping different neighborhood configurations to the same coordinates would introduce systematic errors into the model that cannot be improved by adding more quantum mechanical data. We begin by forming a local atomic density from the neighbors of atom i , as

$$\rho_i(\mathbf{r}) = \delta(\mathbf{r}) + \sum_j \delta(\mathbf{r} - \mathbf{r}_{ij}) f_{\text{cut}}(|\mathbf{r}_{ij}|), \quad (2)$$

where $f_{\text{cut}}(r) = 1/2 + \cos(\pi r/r_{\text{cut}})/2$ is a cutoff function, in which the cutoff radius r_{cut} reflects the spatial scale of the interactions. The choice of cutoff function is somewhat *ad hoc*: any smooth function with compact support could be used.

The local atomic density is invariant to permuting the atoms in the neighborhood. One way to achieve rotational invariance as well would be to expand it in spherical harmonics and a set of radial basis functions and appropriately combine the resulting coefficients, similarly to how the structure factor is computed from Fourier components. However, just as the structure factor (a two-point correlation) is missing all “phase” information (the relative phases of the different plane waves), such a set of spherical invariants would lose a lot of information about the configuration of the neighborhood. In contrast, the bispectrum [9], which is a three-point correlation function, is a much richer system of invariants, and can provide an almost one-to-one representation of the atomic neighborhood.

In our method we first project the atomic density onto the surface of the four-dimensional unit sphere, similarly to how the Riemann sphere is constructed, with the transformation

$$(\phi, \theta, \theta_0) = [\tan^{-1}(y/x), \cos^{-1}(z/|\mathbf{r}|), |\mathbf{r}|/r_0], \quad (3)$$

where $r_0 > r_{\text{cut}}/\pi$. The advantage of this is that the 4D surface contains all the information from the 3D spherical region inside the cutoff, including the radial dimension, and thus 4D spherical harmonics (also called Wigner matrices, $U_{m'm}^j$ [10]) constitute a natural complete basis for the interior of the 3D sphere, without the need for radial basis functions. The projection of the atomic density on the surface of the 4D sphere can therefore be expanded in 4D spherical harmonics using coefficients (dropping the atomic index i for clarity)

$$c_{m'm}^j = \langle U_{m'm}^j | \rho \rangle. \quad (4)$$

The bispectrum built from these coefficients is given by

$$B_{j_1, j_2, j} = \sum_{m'_1, m_1 = -j_1}^{j_1} \sum_{m'_2, m_2 = -j_2}^{j_2} \sum_{m', m = -j}^j (c_{m'm}^j)^* C_{j_1 m_1 j_2 m_2}^{j m'} \times C_{j_1 m'_1 j_2 m'_2}^{j m'} c_{m'_1 m'_2}^{j_1} c_{m'_2 m'_2}^{j_2}, \quad (5)$$

where $C_{j_1 m_1 j_2 m_2}^{j m'}$ are the ordinary Clebsch-Gordan coefficients. The elements of this three-index array, which we will denote by \mathbf{b}_i for atom i , are invariant with respect to permutation of atoms and rotations of 4D space, and hence also 3D space. In practice, we use only a truncated version, with $j, j_1, j_2 \leq J_{\text{max}}$, corresponding to a limit in the spatial resolution with which we describe the atomic neighborhood.

Determining the PES is now reduced to interpolating the atomic energy in the truncated bispectrum space, and for this we use a nonparametric method called Gaussian process (GP) regression [11, 12]. In the GP framework, assuming Gaussian basis functions, the best estimate for the atomic energy function is given by

$$\varepsilon(\mathbf{b}) = \sum_n \alpha_n e^{-(1/2)\Sigma_l [(b_l - b_{n,l})/\theta_l]^2} \equiv \sum_n \alpha_n G(\mathbf{b}, \mathbf{b}_n), \quad (6)$$

where n and l range over the reference configurations and

bispectrum components, respectively, and $\{\theta_{ij}\}$ are (hyper) parameters. The GP is called a nonparametric method because the kernels G are not fixed but centered on the data, and hence, loosely, any continuous function can in principle be obtained from Eq. (6) [13].

The GP differs from a simple radial basis function least-squares fit in the way the coefficients α_n are computed. The covariance, i.e., the measure of similarity, of the reference configurations is defined as

$$C_{nn'} = \delta^2 G(\mathbf{b}, \mathbf{b}') + \sigma^2 \mathbf{I}, \quad (7)$$

where δ and σ are two further hyperparameters and \mathbf{I} is the identity matrix. The interpolation coefficients are then given by

$$\{\alpha_n\} \equiv \boldsymbol{\alpha} = \mathbf{C}^{-1} \mathbf{y}, \quad (8)$$

where $\mathbf{y} = \{y_n\}$ is the set of reference values (quantum mechanical energies). This simple expression for the coefficients is derived in detail in [12]. Thus Eq. (6) gives the atomic energy function in closed form as a function of the quantum mechanical data.

In addition to preserving exact symmetries, another hurdle is that, although we wish to infer the atomic energy function, the data we can collect directly are not values of atomic energies, but total energies of sets of atoms, and forces on atoms, the latter being sums of partial derivatives of neighboring atomic energies [14]. Furthermore, our data will be heavily correlated; e.g., the neighborhoods of atoms in a slightly perturbed ideal crystal are very similar to each other. Both of these problems are solved by applying a sparsification procedure [15], in which a predetermined number (much smaller than the total data size) of “sparse” configurations are chosen randomly from the set of all configurations and the data values \mathbf{y} in Eq. (8) are replaced by linear combinations of all data values. The models in this work used 300 such sparse configurations. The final expression for the model, which we call Gaussian approximation potential (GAP), is derived in [8].

All the DFT data in this work were generated with the CASTEP package [16]. The reference configurations were obtained by randomly displacing the atoms and the lattice vectors from their equilibrium values in 2-, 8-, 16-, and 64-atom unit cells by up to 0.2 Å.

The lower panels of Fig. 1 show the performance of the GAP model for semiconductors in terms of the accuracy of forces for near-bulk configurations. For diamond, the GAP model is shown to improve significantly as the cutoff is increased (there is also systematic improvement as J_{\max} is increased; this is shown in [8]). For all three materials the rms errors in the energy are less than 1 meV/atom. Table I shows the elastic constants. It is remarkable that the existing potentials are not able to reproduce all elastic constants to better than 25% for any setting of their parameters. Figure 2 shows the phonon spectrum for diamond and iron. For diamond, the GAP model shows excellent accuracy at zero temperature over most of the Brillouin zone, with a slight deviation for optical modes in the Λ direction.

TABLE I. Table of relaxed diamond surface energies in J/m² (top) and elastic constants, in units of GPa (bottom).

	DFT			GAP			Brenner			Tersoff (T)		
1×1 unreconstructed	6.41	6.36	4.46	2.85								
2×1 Pandey	4.23	4.40	3.42	4.77								

	C			Si			Ge		
	DFT	GAP	T	DFT	GAP	T	DFT	GAP	T
C_{11}	1118	1081	1072	154	152	143	108	114	138
C_{12}	151	157	108	56	59	75	38	35	44
C_{44}^0	610	608	673	100	101	119	75	75	93
C_{44}	603	601	641	75	69	69	58	54	66

The agreement with experiment is also good for the frequency of the Raman mode as a function of temperature. For iron, the agreement with DFT is even better. We also computed the linear thermal expansion coefficient of diamond, shown in Fig. 3, using two different methods, applicable at low and high temperatures. Our low temperature curve is derived from the phonon spectrum via the quasiharmonic approximation and agrees well with the DFT and experimental results. At higher temperatures higher order anharmonic terms come into play, so we use molecular dynamics (MD) and obtain good agreement with experiment, showing that the GAP model is accurate significantly beyond the small displacements that control phonons.

Finally, we extended the GAP model by including reference configurations generated by random displacements around a diamond vacancy and graphite. Figure 4 shows the energetics of the transition path for a migrating vacancy

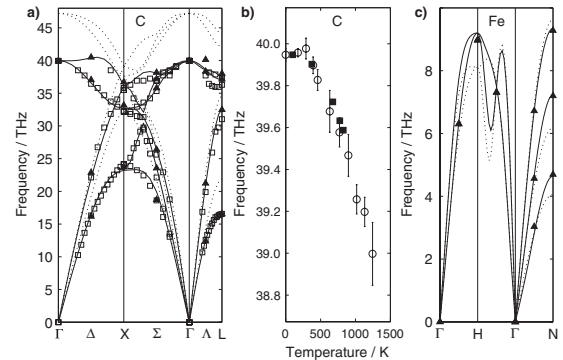


FIG. 2. (a) Phonon dispersion curves for diamond using the GAP model (lines), Tersoff (dotted lines), DFT (triangles), and experiment (squares) [20]. (b) Temperature dependence of the $\Gamma_{25'}$ mode from MD using GAP (circles, 250 atoms, 20 ps) and experiment (squares) [21]. In accordance with common practice [22] the calculated points have been shifted by a constant to agree with experiment at zero temperature to account for the anharmonic effects of zero-point motion and a quantum correction to the kinetic temperature is also applied [23]. (c) Phonon dispersion of iron using GAP (solid), Finnis-Sinclair potential [24] (dotted), and DFT (triangles).




FIG. 3. Linear thermal expansion coefficient of diamond in the GAP model (dashed line) and DFT (dash-dotted line) using the quasiharmonic approximation [25], and derived from MD (216 atoms, 40 ps) with GAP (solid) and the Brenner potential (dotted). Experimental results are shown with squares [26].

in diamond and the transition from rhombohedral graphite to diamond. The agreement with DFT is excellent, demonstrating that we can construct a truly reactive model that describes the sp^2 - sp^3 transition correctly, in contrast to currently used interatomic potentials.

Even for the small systems considered above, the GAP model is orders of magnitude faster than standard plane wave DFT codes, but significantly more expensive than simple analytical potentials. The computational cost is roughly comparable to the cost of numerical bond order potential models [17]. The current implementation of the GAP model takes 0.01 s/atom/time step on a single CPU core. For comparison, a time step of the 216-atom unit cell of Fig. 3 takes 191 s/atom using CASTEP, about 20000 times longer, while the same for iron would take more than 10^6 times longer.

In summary, we have outlined a framework for automatically generating finite range interatomic potential models from quantum-mechanically calculated atomic forces and energies. The models were tested on bulk semiconductors and iron and were found to have remarkable accuracy in matching the *ab initio* potential energy surface at a fraction of the cost, thus demonstrating the fundamental capabilities of the method. Preliminary data for GaN, presented in [8], shows that the extension to multicomponent and charged systems is straightforward by augmenting the local energy with a simple Coulomb term using fixed charges. Our long-term goal is to expand the range of

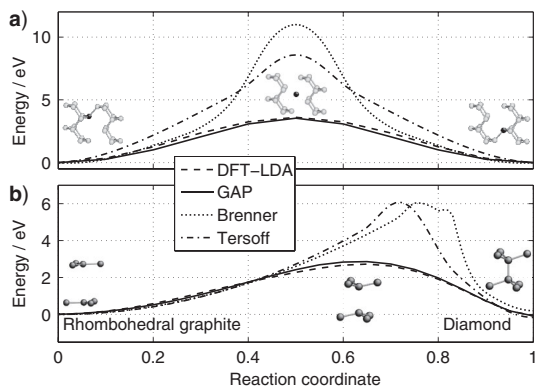


FIG. 4. The energetics of the linear transition path for a migrating vacancy (top) and for the rhombohedral graphite to diamond transformation (bottom).

interpolated configurations and thus create “general” interatomic potentials whose accuracy approaches that of quantum mechanics.

The authors thank Sebastian Ahnert, Noam Bernstein, Zoubin Ghahramani, Edward Snelson, and Carl Rasmussen for discussions. A.P.B. is supported by the EPSRC. Part of the computational work was carried out on the Darwin Supercomputer of the University of Cambridge High Performance Computing Service.

- [1] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover, New York, 1996).
- [2] M. C. Payne *et al.*, Rev. Mod. Phys. **64**, 1045 (1992).
- [3] M. Finnis, *Interatomic Forces in Condensed Matter* (Oxford University Press, Oxford, 2003).
- [4] D. W. Brenner, Phys. Status Solidi B **217**, 23 (2000).
- [5] N. Bernstein, J. R. Kermode, and G. Csányi, Rep. Prog. Phys. **72**, 026501 (2009).
- [6] A. Brown *et al.*, J. Chem. Phys. **119**, 8790 (2003).
- [7] J. Behler and M. Parrinello, Phys. Rev. Lett. **98**, 146401 (2007).
- [8] See supplementary materials at <http://link.aps.org/supplemental/10.1103/PhysRevLett.104.136403> for detailed derivations and results for GaN.
- [9] S. A. Dianat and R. M. Rao, Opt. Eng. (Bellingham, Wash.) **29**, 504 (1990).
- [10] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum* (World Scientific, Singapore, 1987).
- [11] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
- [12] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2003).
- [13] I. Steinwart, J. Mach. Learn. Res. **2**, 67 (2001).
- [14] S. Ahnert, Ph.D. thesis, University of Cambridge, 2005.
- [15] E. Snelson and Z. Ghahramani, in *Advances in Neural Information Processing Systems 18*, edited by Y. Weiss, B. Schölkopf, and J. Platt (MIT Press, Cambridge, MA, 2006), pp. 1257–1264.
- [16] S. J. Clark *et al.*, Z. Kristallogr. **220**, 567 (2005).
- [17] M. Mrovec, D. Nguyen-Manh, D. G. Pettifor, and V. Vitek, Phys. Rev. B **69**, 094115 (2004).
- [18] D. W. Brenner *et al.*, J. Phys. Condens. Matter **14**, 783 (2002).
- [19] J. Tersoff, Phys. Rev. B **39**, 5566 (1989).
- [20] J. L. Warren, J. L. Yarnell, G. Dolling, and R. A. Cowley, Phys. Rev. **158**, 805 (1967).
- [21] M. S. Liu, L. A. Bursill, S. Prawer, and R. Beserman, Phys. Rev. B **61**, 3391 (2000).
- [22] G. Lang *et al.*, Phys. Rev. B **59**, 6182 (1999).
- [23] C. Z. Wang, C. T. Chan, and K. M. Ho, Phys. Rev. B **42**, 11 276 (1990).
- [24] M. W. Finnis and J. E. Sinclair, Philos. Mag. A **50**, 45 (1984).
- [25] P. Pavone *et al.*, Phys. Rev. B **48**, 3156 (1993).
- [26] G. A. Slack and S. F. Bartram, J. Appl. Phys. **46**, 89 (1975).