SCIENTIFIC REPORTS

OPEN

Machine-learned approximations to Density Functional Theory Hamiltonians

Ganesh Hegde & R. Chris Bowen

Received: 22 September 2016

Accepted: 13 January 2017

Published: 15 February 2017

Large scale Density Functional Theory (DFT) based electronic structure calculations are highly time consuming and scale poorly with system size. While semi-empirical approximations to DFT result in a reduction in computational time versus *ab initio* DFT, creating such approximations involves significant manual intervention and is highly inefficient for high-throughput electronic structure screening calculations. In this letter, we propose the use of machine-learning for prediction of DFT Hamiltonians. Using suitable representations of atomic neighborhoods and Kernel Ridge Regression, we show that an accurate and transferable prediction of DFT Hamiltonians for a variety of material environments can be achieved. Electronic structure properties such as ballistic transmission and band structure computed using predicted Hamiltonians compare accurately with their DFT counterparts. The method is independent of the specifics of the DFT basis or material system used and can easily be automated and scaled for predicting Hamiltonians of any material system of interest.

Density Functional Theory (DFT) based electronic structure calculations are a cornerstone of modern computational material science. Despite their popularity and usefulness, DFT methods scale rather poorly with system size. Even the fastest DFT techniques, such as the so-called $O(N)$ approaches based on Local Orbitals spend most of their time in forming the Hamiltonian and solving self-consistently for the ground state electron eigenstates¹. Consequently, while it scales well for small systems such as crystalline solids with small unit cells, scaling the computation of DFT electronic structure for large nanostructures with limited periodicity becomes a challenging task.

Semi-empirical Tight Binding (TB) approximations to DFT perform rather admirably in this regard, reducing the time complexity in forming Hamiltonians by treating Hamiltonian elements as parameters fitted to desired physical properties such as effective mass, band gaps and so on. Typically, TB techniques also simplify the DFT Hamiltonian to near-neighbor (NN) interactions between atoms. This results in extremely sparse Hamiltonians that can be diagonalized very efficiently. Several flavors of these approximations exist, generally optimized to meet the needs of the problem at hand².

While TB calculations have proved immensely useful in modeling the electronic structures of semiconductor and metallic nanostructures, some fundamental problems persist in their use as DFT Hamiltonian approximations. Firstly, while TB Hamiltonians are computationally efficient to solve for, what is often unaccounted for is the time required to create an accurate, physically transparent and transferable TB model for a material system. In our own experience^{3,4}, developing meaningful, physical TB parameter sets has taken time ranging from a few weeks to a few years. Secondly, developing truly transferable, physically valid TB models applicable across materials, geometries and boundary conditions is challenging, owing to physical differences in systems.

There are a variety of reasons for the aforementioned shortcomings. One, it is often the case that a physically transparent model created for a particular material type (say semiconductors with short range, covalent bonding) does not generalize to other material systems (for instance, metals). Secondly, significant portions of the TB parameter-set creation process involve manual intervention. From the functional form of the model used to fit Hamiltonian elements to the choice and tuning of the fitting process, the TB parameterization process often resembles art than an exact science.

This is clearly an undesirable state of affairs, especially in time-sensitive industrial applications where a number of material systems, geometries and boundary conditions need to be evaluated rapidly to test for suitability for an application. To address these issues, we propose a Machine Learning (ML) based method to predict

Advanced Logic Lab, Samsung Semiconductor Inc., Austin, TX 78754, USA. Correspondence and requests for materials should be addressed to G.H. (email: ganesh.h@ssi.samsung.com)

DFT Hamiltonians. In contrast to existing semi-empirical approximations to DFT, our method does not fit parametric models to DFT Hamiltonians or derived electronic structure quantities such as effective masses and band gaps. We instead map atomic environments to real-space Hamiltonian predictions using non-parametric interpolation-based machine learning techniques such as Kernel Ridge Regression (KRR)⁵. The time-consuming direct computation of Hamiltonian integrals and ground state eigenstates in DFT is completely bypassed in favor of simple matrix multiplication reducing computational cost considerably. Through rigorous criteria for selection of algorithm hyper-parameters, the method can be automated and transferability to new material systems can be controlled. The method is independent of the specifics of the material system and depends only on the reference DFT calculations.

The paper is organized as follows. The methods section describes how the problem of predicting DFT Hamiltonians can be mapped on to a ML framework. The results section shows the application of the method proposed in this work to two material systems - an arbitrarily strained Cu system and an arbitrarily strained C (diamond) system. The Cu case study assesses the applicability of the present technique to a relatively simple Cu-DFT system consisting only of rotation invariant *s*-orbitals centered on Cu atoms, while the C case study extends the technique to more realistic systems involving *s* and *p* orbitals and their interactions. A comparison of the advantages and limitations of the method to existing semi-empirical approximations to DFT is included in the discussion section.

Methods

In recent times, several advances have been made in applying a variety of machine learning techniques to computational material science. These include (but are not limited to) the mapping of spatial atomic data to predict total energies^{6–8}, prediction of DFT functionals⁹ and direct computation of electronic properties¹⁰ through machine learning. We are, however, unaware of any attempt to apply ML to the problem of DFT Hamiltonian prediction.

Applying ML to any problem presupposes the existence of data and an unknown, non-trivial map from inputs to outputs. In supervised ML, the dataset consists of inputs/variables (henceforth referred to as ‘features’) and the challenge is to find a map from features to pre-labeled/desired/reference outputs (targets). In our case, the dataset consists of inputs in the form of sets of atomic positions while the targets consist of DFT Hamiltonian elements between orbitals centered on individual atoms. The reference data is obtained by time consuming self-consistent DFT calculations, but the aim is that once a map is obtained from input to reference data, the map is sufficiently accurate and transferable to cases not explicitly included in the dataset.

Real-space DFT Hamiltonians consist of intra-atomic (between orbitals situated on the same atom) and inter-atomic (between orbitals situated on different atoms) matrix elements. Most real-space DFT codes allow direct access to these elements upon the convergence of a ground state DFT calculation. In case of plane-wave DFT codes, a projection of plane-wave based Hamiltonian can be made on to a real space basis. The rest of this paper assumes a real-space Hamiltonian. These real-space DFT Hamiltonian matrix elements then constitute our reference data.

Broadly speaking, intra-atomic and inter-atomic (real-space) DFT Hamiltonian elements depend on

1. The position and charges of the interacting atoms.
2. The positions and charges of atoms that are in the neighborhood of interacting atoms.
3. The basis functions (orbitals) centered on each atom.

(1) and (2) above define a 3D-radially and angularly varying potential (upon DFT calculation convergence) that is highly non-trivial for the simplest of systems that go beyond a single atom. Factor (3) above is usually fully known *a priori*. From this enumeration, one can surmise that obtaining useful features required for ML requires a faithful representation of the positions and charges of interacting atoms and their atomic neighborhood.

A possible starting point in forming features for ML is the collection of atomic positions and charges (atomic numbers). However, operations such as rotation, translation and permutation of atom numbers in an atomic dataset change their positions and their relative ordering completely. In general, Hamiltonian elements are invariant to translation and permutation, but change upon a rotation reference frame (higher spherical harmonics are not rotation invariant). It is therefore important that the features formed reflect these properties. The atomic positions and charges must therefore be transformed to reflect the interactions between orbitals and the atomic environment responsible in forming DFT Hamiltonian interactions.

Seminal work published in the past 10 years has realized the difficult task of invariant representation of atomic neighborhoods possible to the point where these approaches can readily be used in ML for a variety of computational material science applications. These include the Bispectrum formalism of Bartok *et al.*⁶, the Coulomb Matrix formalism (see for instance, Rupp *et al.*⁷ and Montavon *et al.*¹¹), the Symmetry Function formalism of Behler *et al.*⁸, the Partial Radial Distribution Function (PRDF) formalism of Schütt *et al.*¹⁰, the Bag-of-Bonds (BoB) formalism of Hansen *et al.*¹² and the Fourier Series of RDF formalism of von Lilienfeld *et al.*¹³, to name a few. In this work, we have used the bispectrum formalism of Bartok *et al.* to represent the neighborhood of an atom. In this formalism, the 3D radial and angular environment of each atom is converted into vector of bispectrum components of the local neighbor density projected on to a basis of hyperspherical harmonics in four dimensions. While a detailed discussion of the the formalism is beyond the scope of this work, we refer the interested reader to the literature on the topic^{6,14–16}.

For the purpose of this work, it suffices to understand the bispectrum as a unique, vectorized representation of an atoms local environment, so that it can be used as a basis for creating features suitable for machine learning. Figure 1 is a simple example of how the bispectrum distinguishes between different strained environments that a Face Centered Cubic (FCC) Cu atom is exposed to. The various components of the bispectrum vector can reflect subtle variations in neighbor environment, creating a unique map between features and Hamiltonian elements.

Figure 1. Bispectrum measure for a Cu atom in 3 different FCC environments - bulk unstrained, 10% hydrostatic strain and 10% uniaxial strain along [100]. The Bispectrum is a unique measure for the atom in the different environments and is therefore an excellent candidate for the creation of ML features from atomic spatial data.

Our choice of bispectrum is motivated by its rotation, translation and permutation invariant representation of the *atomic environment* in a convenient vectorized form, with an ability to tune the dimensionality of the bispectrum vector by changing the number of angular momentum channels. We hasten to add that the use of the bispectrum should not be construed as distinctly advantageous by the reader and alternate representations of the atoms environment that share these features may be considered in place of the bispectrum without affecting the applicability of the present method.

s-orbitals and their interactions with other *s*-orbitals are rotation, translation and permutation invariant. For the Cu-case study where a single *s*-orbital centered on Cu atoms forms the basis set, our input features can then be simply formed from the bispectra of the atoms taking part in the interaction without explicit information about the reference frame or its rotation. For intra-atomic elements in the Cu-case study, the feature vector we use is simply the bispectrum of the respective atoms. For inter-atomic Hamiltonian elements in this case study, a candidate feature vector is the bispectrum of the two atoms taking part in the interaction. Since 3D periodic material systems contain an infinite number of atoms having the same neighborhood, we divide the bispectra of the two atoms taking part in the Hamiltonian interaction by their inter-atomic distance and append the interatomic distance to form the feature vector. We found that adding some measure of the interatomic distance explicitly in the feature vector in addition to its implicit inclusion by (modifying the bispectrum) systematically improved predictive performance of our method. For instance we found that the RMSE of the inter-atomic interactions reduced by approximately 2 meV with the inclusion of interatomic distance. We note, however, that we did not systematically investigate the performance of a distance metric being explicitly included in the feature vector versus, say, inverse distance or other powers of interatomic distance. Figure 2 in the methods section shows the feature formation process for the Cu-case study in greater detail.

For the C (diamond) - case study, the simplest DFT basis set that captures electronic interactions accurately is one comprised of a single *s* orbital and 4 *p* orbitals. While interactions between *s* orbitals are rotation invariant as mentioned previously, interactions between dissimilar *s* and *p* orbitals and between dissimilar *p* orbitals on different atoms in general are not rotation invariant. In addition to capturing the atomic environment, which may yet be invariant to rotations, one therefore needs to capture angular information as well. For interatomic interactions, the direction cosines (*l*, *m*, *n*) of the bond vector (with respect to the Cartesian axes/laboratory reference frame) between two atoms taking part in the Hamiltonian interaction were appended to the feature vector used in the Cu-case study. Introduction of this angular information simply captures the necessary angular information for interatomic elements.

The orbitals involved in intra-atomic interactions belong to the same atom, while the potential contribution to the Hamiltonian element is from this atom and other neighboring atoms. Rotation of the reference frame or addition of atoms in the neighborhood leaves the relative angular information of the orbitals situated on the same atom unchanged i.e. still zero. The angular information of the neighborhood, however, changes and with this rotation, the off-diagonal intra-atomic Hamiltonian elements change as well. Unlike inter-atomic interactions, however, we cannot solve this problem by simply appending neighbor direction cosines, since the interaction captured is between two orbitals on the same atom. This is addressed by appealing to the physics of the problem at hand. The neighbors that exert the most influence on the intra-atomic interaction are the nearest neighbors. We therefore sum the direction cosines with respect to each Cartesian axis from each neighbor and append this to the bispectrum. For instance, a C atom in diamond-like configurations has 4 nearest neighbors, with direction cosines l_1, l_2, l_3 and l_4 with respect to the *x* Cartesian axis. For unstrained diamond, the sum of these direction cosines is zero as is the sum of direction cosines with respect to *y* and *z* axes. This captures very well the actual DFT calculation for the intra-atomic block in diamond, where off-diagonal intra-atomic interactions are zero. When symmetry is broken by straining uni-axially for instance, the sum of direction cosines becomes non-zero. This is reflected

Figure 2. Schematic representation of the reference dataset creation process for the Cu-case study. First the atomic structure data in the form of atom types, charges and positions are converted to corresponding bispectra. The bispectra of different atoms are then combined to form features for learning. For intra-atomic Hamiltonian elements (E), the bispectra of individual atoms constitute the feature vector. For inter-atomic Hamiltonian elements (V), the bispectra are divided by the inter-atomic distances (R) to form the feature vector. The feature vectors along with their corresponding reference output data constitute the reference dataset.

in the corresponding Hamiltonian where the off-diagonal intra-atomic elements become non-zero. This scheme outlined above is one among a number of such schemes that can be devised to incorporate angular interactions.

In the Slater-Koster scheme used in classical Tight-Binding¹⁷, introducing angular information in the form of direction cosines can be shown to represent the angular part of all inter-atomic Hamiltonian interactions, while canonical interatomic interactions between orbitals - $V_{ss\sigma}$, $V_{sp\sigma}$, $V_{pp\pi}$, $V_{pp\sigma}$ - can be obtained by aligning the bond-axis to the laboratory/Cartesian z-axis reference. Generally speaking, however, interatomic interactions go beyond the simple-two center scheme and the Slater Koster scheme applied fails to account for more complicated 'three-center' interactions. This is true for intra-atomic interactions as well, where one needs to go beyond the two-center approximation to capture non-zero off-diagonal intra-atomic Hamiltonian elements. Since the aim in machine learning is typically to let the data 'find' a reliable map upon supplying features sufficient to represent the outputs meaningfully, we have simply chosen to incorporate angular information in the form of direction cosines and their sums instead of imposing a model on the underlying interactions.

Having formed the features thus, the non-linear map from the features to the Hamiltonian elements can be learned using a ML technique called Kernel Ridge Regression (KRR). In KRR, the predicted output for a given feature vector is expressed as a combination of similarity measures between the feature vector (for which an output is to be predicted) and all other feature vectors in the dataset. Using the notation from Rupp¹⁸, for a feature vector $\tilde{x}$, the predicted output $f(\tilde{x})$ is given as

$$f(\tilde{x}) = \sum_{i=1}^n \alpha_i k(x_i, \tilde{x}) \quad (1)$$

Here, n is the number of training examples, $k(x_i, \tilde{x})$ is the kernel measure between feature vectors x_i and $\tilde{x}$ and α_i is the corresponding weight in the sum. The kernel measure $k(x_i, \tilde{x})$ allows a comparison of similarity (or conversely, distance) between feature vectors⁵. We use a Gaussian kernel defined as follows

$$k(x_i, \tilde{x}) = \exp(-\gamma \|x_i - \tilde{x}\|_2^2) \quad (2)$$

γ is an algorithm hyper-parameter that decides the extent (in feature space) to which the kernel operates. The second term in the exponent is the squared 2-norm of the difference vector between the two feature vectors x_i and $\tilde{x}$. The similarity kernel returns a value of 1 when the feature vectors are identical, indicating maximum similarity. By changing the value of γ one can tune the extent to which the kernel returns similarity values. For γ values close to zero, only those feature vectors that are closely matched will return a similarity measure approaching 1, whereas feature vectors that are dissimilar will return values approaching zero. For larger values of γ , a finite similarity value is returned by the kernel for feature vectors that are unequal.

Equation 1 above thus interpolates between existing feature vectors in the dataset and predicts a value based on the similarity of the candidate feature vector to those in the dataset. The weights α are learned by computing the similarity measure between all feature vectors in the training dataset.

$$\alpha = (K + \lambda I)^{-1} y \quad (3)$$

Figure 3. Grid of algorithm hyper-parameters γ and λ versus Root Mean Square Error (RMSE) scores for the V dataset for the C-diamond case study. A large number of possible combinations of hyper-parameters result in a similar range of RMSE and can be used for training weight α . We used the combination that provided the lowest RMSE across both case studies.

where the kernel matrix K is obtained by computing similarity measures using equation 2 between each feature vector in the training dataset, y is the vector of training reference data. λ is a regularization hyper-parameter that controls the smoothness of the interpolator and prevents over-fitting. I is the identity matrix.

Equations 1, 2 and 3 together constitute the equations defining the KRR learning algorithm. Given a training dataset X (with n feature vectors x_i and i from 1 through n) and y (with reference outputs y_i for each feature vector in the training dataset), these equations can be used to train the set of weights α , obtain a similarity measure K for the dataset, and predict the output value for a candidate feature vector $\tilde{x}$ for a given value of the algorithm hyper-parameters λ and γ .

Having shown how a DFT Hamiltonian problem can be mapped on to a machine learning problem, we now summarize the requirements from a procedural perspective. First, n reference self-consistent DFT calculations on the material system of interest are performed. For each DFT calculation, DFT Hamiltonian elements are extracted to form the reference outputs y . For each reference output y_i , a feature vector x_i is extracted by converting the atomic positions, charges and neighborhoods of the atoms involved in the formation of matrix elements y_i to a corresponding invariant bispectrum measure using the aforementioned procedure. The reference dataset (X, y) thus formed is partitioned into training (X_{train}, y_{train}) and test datasets (X_{test}, y_{test}) . Figure 2 is a schematic representation of the creation of the procedure used to create the reference datasets.

The training dataset is used to learn the weights w that map feature vectors X_{train} to outputs y_{train} using equation 2 and 3 above for the set of hyper-parameters λ and γ . Equation 1 can then be used to predict Hamiltonian elements. The quality of the fit can be assessed by comparing predicted values $f(\tilde{X}_{test})$ for features in the test dataset with the reference values y_{test} .

The selection of algorithm hyperparameters λ and γ is done by creating a grid of λ and γ values and evaluating a goodness-of-fit metric for each combination of hyperparameters in the grid. A number of statistical metrics such as the R^2 score, the root mean square (RMS) error, the absolute error can be used. The hyperparameter combination that results in the best goodness-of-fit metric can be used for training and prediction. Using this technique, algorithm hyper-parameters can be rigorously and automatically chosen instead of basing the selection on trial-and-error. Figure 3 is an example of a goodness-of-fit metric plotted versus a grid of algorithm hyperparameters. The dataset generated initially can be augmented continuously and the algorithm weights and hyper-parameters can easily be re-learned for maximum transferability. In addition, a significant portion of the prediction procedure - learning weights, hyper-parameters and predicting Hamiltonian elements - can be parallelized. For large datasets, efficient algorithms such as gradient descent can be used instead of the linear equations above to optimize algorithm weights⁵. The accuracy and transferability of the technique can be further improved by techniques such as K -fold cross validation, where the training and validation is repeatedly performed on smaller partitions of the dataset^{5,18} and testing is performed on the test dataset that does not enter the fitting procedure. We used $K = 5$ -fold cross validation for all the results in this work.

Our ML technique of choice in this work was KRR with Gaussian kernels. We mention, however, that our choice of kernel is not unique. For instance, Laplacian kernels applied to the study of molecular energies¹⁸ appear to have improved performance in comparison to Gaussian kernels. Alternative machine learning algorithms such as Gaussian Process Regression (GPR)¹⁹ and Support Vector Regression (SVR)²⁰ can perform the same task with similar accuracy and ease of implementation²¹.

Since the feature vectors created incorporate charge and distance information, the method is independent of the specifics of the material system used and the same algorithm can be applied without change to a variety of material systems. The material system and DFT specifics-agnostic nature of the method makes it suitable for high-throughput electronic-structure screening calculations where a number of machines can be learned rapidly

Figure 4. Predicted versus reference intra-atomic interaction energies (left) and inter-atomic interaction energies (right) for the single s -orbital Cu system studied. The $x = y$ line is included as an indicator for the quality of fit - when predictions match reference/desired data perfectly, they lie on this line.

for multiple materials and Hamiltonian predictions can be made to account for changes in electronic structure versus orientation, periodicity, disorder and so on.

Results

As mentioned previously, we demonstrate the soundness of the method with the help of two case studies. The first involves the prediction of Hamiltonian elements in hydrostatically and uniaxially strained FCC Cu. For simplicity and ease of demonstration, we use a compensated single s -orbital basis for Cu^{22} for DFT calculations. The second case study, involving multi-orbital systems involved the uniaxial and hydrostatically strained diamond phase of Carbon.

Single s -orbital strained Cu system. For the single s -orbital basis used in the Cu-case study, the DFT Hamiltonian consists of only two types of elements

1. E_s - The intra-atomic or on-site energy element for the Cu s orbital. We henceforth refer to this element as E . E captures the effect of the external environment (potential created by the neighborhood of atoms) on a given Cu atom.
2. V_{ss} - The inter-atomic or off-site energy element that captures the strength of the interaction between s -orbitals situated on interacting atoms mediated by the potential created by the neighborhood of interacting atoms. We henceforth refer to this element as V .

To create the reference dataset, we performed self-consistent DFT calculations for a total of 200 uniaxially and hydrostatically strained bulk FCC Cu systems. This results in a reference/labeled dataset with 27000 intra and inter-atomic Hamiltonian elements for the strained Cu system. The details of the DFT calculations for Cu such as the basis set, the compensation charge, the DFT functional, the k -point grid density are outlined in work we recently published²². The specifics of the DFT technique used are unimportant for this work, since the ML technique is independent of the type and details of DFT calculation. One only needs to ensure that a consistent k -point grid density and DFT functional is used across all reference calculations so that the calculations are consistent across boundary conditions.

From the converged DFT calculations, we extract E values for each atom in the calculation and V values for each (atom, neighbor) combination to form reference output data y_E and y_V . The E reference data set for the Cu case-study consisted of 500 intra-atomic Hamiltonian elements. The feature vector for E prediction for an atom is a row vector consisting of 69 bispectrum coefficients for that atom. The reference feature matrix X_E is thus a $N \times 69$ matrix where $N = 500$ is the total number of atoms across all reference calculations for the Cu case study. The V reference dataset consisted of 26500 Hamiltonian elements. The V dataset is significantly larger than the E dataset since it captures interactions between each of the atoms in the E dataset and *all* of its neighbors. The spatial extent of the s -orbital used in the DFT calculations is up to approximately 8 Angstrom. Consequently, each atom in the E dataset has (averaging over all strain configurations), approximately 133 neighbors for Cu. The feature vector for V prediction consists of 69 bispectrum coefficients for the first and second atoms each divided by the interatomic distance. The interatomic distance is also appended at the end of this vector for a total dimensionality of 139 for V features in the dataset. The reference feature matrix X_V is thus a 26500×139 matrix.

Machines for E and V prediction were then trained using the procedure outlined previously. We use a randomized training:test split of 80:20 percent for E and V training and prediction. We use the Root Mean Square Error (RMSE) between reference and predicted test data as a goodness-of-fit metric to test the quality of our predictions.

We achieved a RMSE of approximately 0.3 meV for the E dataset and an RMSE value of approximately 6 meV for the V dataset in this case-study. Figure 4 shows a comparison of predicted versus reference E and V values for the training and test dataset. The $x = y$ line is also included as an indication of the quality of the prediction. A perfect prediction matching reference data lies on this line. It can be seen that the E dataset is trained and

predicted with high fidelity. The V dataset also has predictions that mostly lie on or near the $x = y$ line, with some interactions close to zero being predicted away from the line.

sp³-orbital strained C (diamond) system. For a 4 orbital sp³ basis used in the C-case study, the DFT Hamiltonian consists of 3 categories of elements.

1. E_s, E_p - The intra-atomic or on-site energy element for the C s and p orbitals respectively that capture effect of the neighborhood of atoms on a given orbital at an atomic site.
2. $E_{s,p_{x,y,z}}, E_{p_x,p_{y,z}}$ etc. - Off-diagonal intra-atomic elements that capture interactions between s and p orbitals on the same atom mediated by the potential created by the neighborhood of atoms.
3. $V_{s,s}, V_{s,p_{x,y,z}}, V_{p_{x,y,z},p_{x,y,z}}$ - The inter-atomic Hamiltonian element that captures the strength of the interaction between orbitals situated on different atoms mediated by the potential created by the neighborhood of interacting atoms.

For the multi-orbital basis, the intra-atomic elements are collectively referred to as E, while the inter-atomic elements are collectively referred to as V. We note again that since the two-center approximation does not generally hold in DFT (or in realistically strained systems), we cannot resolve interactions between orbitals in the form of $V_{sp\sigma}, V_{pp\pi}, V_{pp\sigma}$ and other such canonical interactions and we must therefore treat each DFT Hamiltonian element category as a distinct reference output.

To create the reference dataset, we performed self-consistent DFT calculations for a total of 40 strained C systems. The unit cells ranged from single atom FCC unit cells strained arbitrarily to within 4% of the lattice parameter to larger unitcells oriented along [100], [110] and [111] transport orientations. From the converged DFT calculations, we extract E values for each (atom, orbital-orbital) combination in the calculation and V values for each (atom, orbital, neighbor) combination to form reference output data y_E and y_V . This results in a reference/ labeled dataset with approximately 45000 Hamiltonian elements for the C system. A single-zeta numerical basis set consisting of a single s and 3 p orbitals was used as the basis set, with norm-conserving pseudopotentials within the Local Density Approximation (LDA) in DFT²³. Our observation was that the same k-grid density used for Cu was sufficient to ensure converged self-consistent total energy in the C-calculations as well. We use LDA in spite of the well known band-gap underestimation problem for semiconductors because our emphasis in this work is not how well DFT calculations represent the real band structure of C-diamond. Rather, our focus is on how well the method we propose predicts DFT Hamiltonian elements and derived quantities. Being DFT-specific agnostic, the technique is rather independent of the details of the DFT calculation. As in the case study for strained Cu, one only needs to ensure that a consistent k-point grid density and DFT functional is used across all reference calculations so that the calculations are consistent across boundary conditions.

As in the case of the Cu-case study, from the converged DFT calculations, we extract E values for each [atom, orbital type] combination in the calculation and V values for each [atom, orbital type, neighbor] combination to form reference output data y_E and y_V . The E reference data set for the C case-study consisted of intra-atomic Hamiltonian elements corresponding to a total of 800 atoms in the dataset. Each distinct interaction was represented by a separate column in the reference/output data-matrix. For instance $E_{s,s}, E_{s,p_x}$ and so on are represented by a distinct column in the output data matrix. The feature vector for E prediction for an atom is a row vector consisting of 69 bispectrum coefficients for that atom to which the sum of nearest neighbor direction cosines $\sum l_i, \sum m_i, \sum n_i$ (where index i is over the nearest neighbors) are appended. The reference feature matrix X_E is thus a $N \times 72$ matrix where $N = 800$ is the total number of atoms across all reference calculations for the C-case study. The E output/DFT reference data-matrix is a 800×10 matrix.

The V reference dataset consisted of 44200 Hamiltonian elements. Like in the case of Cu, the V dataset is significantly larger than the E dataset. The feature vector for V prediction consists of 69 bispectrum coefficients for the first and second atoms each divided by the interatomic distance. The 3 direction cosines for the bond corresponding to each atom, neighbor combination are appended to this vector followed by the interatomic distance at the end for a total dimensionality of 142 for V features in the dataset. The reference feature matrix X_V is thus a 44200×142 matrix. The output data matrix is a 44200×10 matrix, where each distinct interatomic interaction is captured in a distinct column in the output dataset.

As in the case of the Cu case study, machines for E and V prediction were then trained using the procedure outlined previously. We also use a randomized training:test split of 80:20 percent for E and V training and prediction. Similar to the Cu-case study, we use the RMSE as a measure of goodness-of-fit. We achieved a RMSE of 1 meV for the E dataset and an RMSE of 18 meV for the V dataset. Figure 4 shows a comparison of predicted versus reference E and V values for the training and test dataset in the C-diamond case-study.

While accurate and computationally efficient prediction of Hamiltonians is a goal of this work, our proposal can be judged as useful only if electronic structure quantities derived from predicted Hamiltonians compare accurately to their *ab initio* DFT calculated counterparts. Since these quantities are not explicitly included in the fitting process, such a comparison also serves as an additional validation of the accuracy of the predicted Hamiltonians. We use band structure and ballistic electronic transmission as a figure-of-merit to compare quantities calculated from predicted Hamiltonians and DFT Hamiltonians computed *ab initio*. For Cu, in the spirit of our previous investigations on the topic^{24,25}, we use ballistic transmission as a figure of merit. While ballistic transmission for perfectly periodic structures can be derived from band structure, the applications we cited above make no assumption for periodicity. For C-diamond, which is a semiconductor, band structure is a more appropriate figure of merit. This is because interesting spectral features such as the minimum energy band gap and effective mass, crucial for semiconductors, can be derived directly from band structures.

Figure 5. Examples of ballistic Transmission computed using the Hamiltonians predicted by ML versus the same calculation done *ab initio* in DFT for the single *s*-orbital Cu system studied. From left to right these are (1) transverse-momentum resolved transmission in a bulk unstrained FCC Cu unit cell oriented along [100] (2) transverse momentum resolved transmission in a [110] 2% tensile uniaxially strained Cu unit cell oriented along [110]. (3) Transverse momentum resolved transmission in a 2% compressive hydrostatically strained Cu unit cell oriented along [111].

Figure 6. Predicted versus reference transverse-momentum Averaged Transmission calculated for DFT Hamiltonian elements in the test dataset.

Figure 5 is an example of such a comparison. The figure shows transverse k -resolved ballistic transmission for bulk-unstrained FCC Cu oriented along [100] and for a [110] unit cell uniaxially strained by 2% along transport direction. We note here that in spite of not having directly included k -resolved transmission information in the training process, the DFT-reference and ML-predicted transmission spectra compare quite well qualitatively. A (unitless) transmission RMSE of 0.1 averaged across all k -grid values (21×21 k grid per structure) for each structure in the dataset was achieved. Figure 6 shows a comparison of the transmission averaged over the transverse Brillouin zone. We computed a transmission RMSE of 0.08 between the reference and predicted *average* transmissions, across a transmission range of 4 (maximum transmission = 4, minimum transmission = 0) for the structures studied indicating excellent agreement for the predictions.

The band structure of unstrained C-diamond is predicted and contrasted with its DFT-calculated counterpart in Fig. 7. We note here that in the case of semiconductor band structure modeling using SETB, it is customary to target a limited set of bands that are useful from a conductance perspective. Only those bands that are close to within 25 meV of the Fermi Level are interesting from this point of view. Even so, since the Fermi Level can be modulated by *p* and *n* type doping, it is necessary to predict the band structure of the highest valence and lowest

Figure 7. From left to right - Predicted versus reference intra-atomic interaction energies (left), inter-atomic interaction energies (center) for the sp^3 -orbital C-diamond system studied. The $x = y$ line is included as an indicator for the quality of fit - when predictions match reference/desired data perfectly, they lie on this line. (right) The band structure of unstrained diamond derived from a prediction of the Hamiltonian using this technique.

conduction band with high degrees of accuracy. We obtained a band structure RMSE of 7 meV for the lowest conduction and highest valence band. If we average across all bands for all test structures we find that our RMSE goes up to 0.1 eV. Data mining our predicted Hamiltonian values reveals that this arises chiefly due to $p - p$ interactions that are predicted with a higher RMSE of 50 meV (since $p - p$ interactions contribute significantly in bands well above the Fermi Level in undoped C-diamond). A visual inspection of the C-diamond band structure in 7 confirms a significantly accurate prediction of the valence bands in C while the conduction bands seem to be predicted accurately for some regions in k-space, while they deviate from DFT predictions considerably in other regions.

Discussion

The results in previous sections indicate that learning and predicting DFT Hamiltonians is not only feasible, but can also be done with high accuracy of predictions provided meaningful features linking structure to Hamiltonian elements can be created. It is useful to contrast the method proposed in this paper to other DFT-approximations that exist in the literature in order to understand its relative advantages and limitations. Existing approximations to DFT Hamiltonians can be broadly classified into two main types. In the first, the Hamiltonian elements are computed by fitting a model of intra and inter-atomic interactions to eigenspectra and/or wavefunctions^{4,26,27}. These techniques create a mathematical map from inputs (atom types, their positions and derived TB Hamiltonian matrix elements) to DFT derived-outputs (wavefunctions, eigenspectra and derived quantities such as band gaps and effective masses). In the second, instead of fitting a model to DFT derived outputs, one alternately maps basis functions and Hamiltonian elements (inputs) to DFT Hamiltonian matrix elements (outputs). If a highly accurate map is found, the derived quantities reflect a similar accuracy. Explicit projection of DFT basis sets and Hamiltonian elements to lower order TB Hamiltonians has been discussed in detail in several publications^{28–31}.

In contrast to the latter set of techniques, we do not form an explicit mathematical map to DFT Hamiltonians. Our method instead interpolates between feature-vectors that map non-linearly to Hamiltonian elements. Instead of defining the map using arbitrary mathematical models that link features to outputs, the use of the kernel-trick in KRR allows defining the map in a high-dimensional feature space without needing to make the map explicit. Since the map changes implicitly for different material systems, we can use the method as is without needing to account for explicit changes in the physics of the material system. In other words, a change in material systems necessitates only re-training of weights for a new dataset that may be generated for this material system. This in contrast, to say, TB, where the use of explicit functional forms relating the parameterized Hamiltonian to its environment requires a continuous assessment of the validity of the model across material systems. Model transferability from one material system to the other is often poor, due to which a wide variety of semi-empirical models exists in the literature (see Goringe *et al.*² for an example of the wide variety of TB models created for different material systems).

No additional assumption regarding the orthogonality of the basis set needs to be made when using Hamiltonian elements predicted by this method. Given that the overlap between numerical basis sets used in real-space DFT codes can completely be defined by doing a dimer calculation for varying radial distance, the overlap integral component of the generalized eigenvalue problem is considered as solved. One can simply pre-compute overlap integrals versus radial distance and use this calculation as a look-up table when needed.

Outside of the decision regarding the best set of physically transparent features to use for the problem at hand, our method obviates the need to fit eigenvalues and other derived features that characterizes almost all semi-empirical approximations to DFT. It is also evident that the formation of the Kernel, training of weights and the selection of hyper-parameters are entirely automated. The creation of training data by can be done by scripting self-consistent DFT calculation runs. Even though other semi-empirical approaches save considerable computational time by avoiding explicit computation to self-consistency, the amount of time taken to develop a

particular approximation is never accounted for. In the past, this has been a significant roadblock in our own use of TB approximations to DFT for high-throughput electronic structure screening calculations.

The KRR technique automatically leads to the ability to incorporate multiple outputs/targets within the same procedure by simply adding columns to the output dataset. A new set of weights α for each output of interest is generated in the training process. This is quite convenient since multi-orbital interactions can now be extended very simply in case one decides to incorporate a larger basis set with an increased number of orbitals.

Accuracy of the technique in describing specific spectral features can be systematically improved by targeting Hamiltonian elements that govern those spectral features. For instance, in the C-diamond case study, upper conduction bands were not predicted as accurately as valence bands and lower conduction bands. This was attributed to the lower accuracy of the $p-p$ interaction prediction. With the use of weights in the KRR technique, $p-p$ interactions can be targeted for systematic improvement. This, however, comes at the cost of arbitrariness in deciding weights that we prefer avoiding. It appears that alternate ways to systematically improve prediction performance that are grounded in ML theory may be needed for this purpose. For instance, an alternate feature set that leads to a more accurate map for $p-p$ interactions may be needed. Even so, we note that in our experience, generating a model in SETB for arbitrary strain in semiconductors is a venture that results in a development time running into several months at minimum.

Some of the advantages of our method versus alternate DFT-approximation techniques also result in its most important limitation. Creating an explicit model for Hamiltonian elements allows a deeper understanding of the physics of the system, even if the model is not transferable across material systems. For instance, the TB models of Wang *et al.*³² and Pettifor *et al.*³³, relate changes in the atomic environment to Hamiltonian parameters using an intuitive model of how bonding with neighbors affects inter-atomic interactions. Since our technique is non-parametric, this intuitive understanding is lost. In this regard, our opinion is that the opportunity-cost of intuitive understanding when using explicit models is more than compensated for by the ease of implementation, the degree of automation and transferability our method provides.

An alternative route to learning electronic properties directly from structural information has also been adopted in recent literature. For instance, the work of Lopez-Bezanilla³⁴ and Schütt *et al.*¹⁰ shows that the electronic transmission and density of states can be linked to structural information directly. While this approach is immensely powerful and efficient for specific properties, having access to an accurate Hamiltonian prediction enables the computation of all of all such derived quantities simultaneously instead of having to train multiple machines on multiple datasets with different reference electronic properties.

The advantage of supervised machine learning techniques in making sense of existing, labeled data is also an important limitation in some sense. It is clear from the KRR formalism presented above that the technique is interpolation-based. In other words, the quality of the prediction is dependent on having data representations in the training data set that are somewhat similar to the candidate for prediction. In certain cases, this is a limitation that can be ignored - for instance, the space of known and achievable strains in Si used in logic technology lead to a rather small number of interesting cases of strained Si within the space of all possible strains and deformations of FCC-Si. In such a case it is very easy to do a dense sampling of strained Si candidates, so that predictions on this system can be interpolated with high accuracy. Predictions of the band structure of silicene from a data set trained on strained FCC-Si, however, are not feasible in KRR (and possibly in other ML techniques) due to their supervised nature, where some labeled data is required in advance of the prediction. This necessitates enlarging the training dataset to incorporate at least some sense of the physical interactions that may be encountered in a target application.

Extensions of the technique to multi-component/multi-element systems involves two distinct steps. The first is the computation of a bispectrum/alternate-feature vector for a multi-component system. The second is extending the number of interactions that need to be mapped from DFT using machine learning. As a case in point, we consider the semiconductor InAs. Instead of a single $s-p$ interaction we included in the output reference dataset, now, two such interactions need to be included, since the $s-p$ interactions involving an In s orbital and an As p orbital is quite different than those involving As s and In p orbitals. Apart from this modification, no additional changes are needed in the technique when applying it to multi-component systems.

In conclusion, we have proposed a technique based on Machine Learning for automated learning and prediction of DFT Hamiltonians. The method is based on a non-parametric, non-linear mapping of atomic bispectrum features to DFT Hamiltonian elements using Kernel Ridge Regression. It is transferable across material systems and is independent of the specifics of the DFT functional used. We have demonstrated that the method can provide accurate prediction of DFT Hamiltonians and that the derived electronic structure quantities compare accurately to their DFT calculated counterparts. We believe that the method provides a promising starting point for highly-automated learning and prediction of DFT Hamiltonian elements towards application in high-throughput electronic structure calculations.

References

1. Bowler, D. & Miyazaki, T. Methods in electronic structure calculations. *Reports on Progress in Physics* **75**, 036503 (2012).
2. Goringe, C., Bowler, D. *et al.* Tight-binding modelling of materials. *Reports on Progress in Physics* **60**, 1447 (1997).
3. Klimeck, G. *et al.* Si tight-binding parameters from genetic algorithm fitting. *Superlattices and Microstructures* **27**, 77–88 (2000).
4. Hegde, G., Povolotskyi, M., Kubis, T., Boykin, T. & Klimeck, G. An environment-dependent semi-empirical tight binding model suitable for electron transport in bulk metals, metal alloys, metallic interfaces, and metallic nanostructures. i. model and validation. *Journal of Applied Physics* **115** (2014).
5. Friedman, J., Hastie, T. & Tibshirani, R. *The elements of statistical learning* vol. 1 (Springer series in statistics Springer, Berlin, 2001).
6. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Physical review letters* **104**, 136403 (2010).
7. Rupp, M., Tkatchenko, A., Müller, K.-R. & Von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Physical review letters* **108**, 058301 (2012).

8. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical review letters* **98**, 146401 (2007).
9. Snyder, J. C., Rupp, M., Hansen, K., Müller, K.-R. & Burke, K. Finding density functionals with machine learning. *Physical review letters* **108**, 253002 (2012).
10. Schütt, K. *et al.* How to represent crystal structures for machine learning: Towards fast prediction of electronic properties. *Physical Review B* **89**, 205118 (2014).
11. Montavon, G. *et al.* Learning invariant representations of molecules for atomization energy prediction. In *Advances in Neural Information Processing Systems*, 440–448 (2012).
12. Hansen, K. *et al.* Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *The journal of physical chemistry letters* **6**, 2326–2331 (2015).
13. von Lilienfeld, O. A., Ramakrishnan, R., Rupp, M. & Knoll, A. Fourier series of atomic radial distribution functions: A molecular fingerprint for machine learning models of quantum chemical properties. *International Journal of Quantum Chemistry* **115**, 1084–1093 (2015).
14. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Physical Review B* **87**, 184115 (2013).
15. Thompson, A. P., Swiler, L. P., Trott, C. R., Foiles, S. M. & Tucker, G. J. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *Journal of Computational Physics* **285**, 316–330 (2015).
16. Bartók, A. P. & Csányi, G. Gaussian approximation potentials: A brief tutorial introduction. *International Journal of Quantum Chemistry* **115**, 1051–1057 (2015).
17. Podolskiy, A. & Vogl, P. Compact expression for the angular dependence of tight-binding hamiltonian matrix elements. *Physical Review B* **69**, 233101 (2004).
18. Rupp, M. Machine learning for quantum mechanics in a nutshell. *International Journal of Quantum Chemistry* **115**, 1058–1073 (2015).
19. Rasmussen, C. E. Gaussian processes for machine learning (2006).
20. Smola, A. J. & Schölkopf, B. A tutorial on support vector regression. *Statistics and computing* **14**, 199–222 (2004).
21. Abu-Mostafa, Y. S., Magdon-Ismael, M. & Lin, H.-T. *Learning from data* vol. 4 (AMLBook Singapore, 2012).
22. Hegde, G. & Bowen, R. C. On the feasibility of ab initio electronic structure calculations for Cu using a single s orbital basis. *AIP Advances* **5**, 107142 (2015).
23. Manual, A. T. Atk version 2015.1. *QuantumWise A/S* (www.quantumwise.com).
24. Hegde, G., Bowen, R. & Rodder, M. S. Lower limits of line resistance in nanocrystalline back end of line Cu interconnects. *Applied Physics Letters* **109** (2016).
25. Hegde, G., Bowen, R. C. & Rodder, M. S. Is electron transport in nanocrystalline Cu interconnects surface dominated or grain boundary dominated? In *Interconnect Technology Conference/Advanced Metallization Conference (IITC/AMC)*, 2016 IEEE International 114–116 (IEEE, 2016).
26. Tan, Y. P., Povolotskyi, M., Kubis, T., Boykin, T. B. & Klimeck, G. Tight-binding analysis of Si and GaAs ultrathin bodies with subatomic wave-function resolution. *Physical Review B* **92**, 085301 (2015).
27. Tan, Y., Povolotskyi, M., Kubis, T., Boykin, T. B. & Klimeck, G. Transferable tight-binding model for strained group IV and III-V materials and heterostructures. *Physical Review B* **94**, 045311 (2016).
28. Urban, A., Reese, M., Mrovec, M., Elsässer, C. & Meyer, B. Parameterization of tight-binding models from density functional theory calculations. *Physical Review B* **84**, 155119 (2011).
29. Elstner, M. *et al.* Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties. *Physical Review B* **58**, 7260 (1998).
30. Wang, C.-Z. *et al.* Tight-binding hamiltonian from first-principles calculations. In *Scientific Modeling and Simulations*, 81–95 (Springer, 2008).
31. Qian, X. *et al.* Quasiatomic orbitals for ab initio tight-binding analysis. *Physical Review B* **78**, 245112 (2008).
32. Wang, C. *et al.* Environment-dependent tight-binding potential model. In *MRS Proceedings*, vol. 491, 211 (Cambridge Univ Press, 1997).
33. Pettifor, D. New many-body potential for the bond order. *Physical review letters* **63**, 2480 (1989).
34. Lopez-Bezanilla, A. & von Lilienfeld, O. A. Modeling electronic quantum transport with machine learning. *Phys. Rev. B* **89**, 235411 (2014).

Acknowledgements

We acknowledge useful discussions with Borna Obradovic and Ryan Hatcher. We thank Mark Rodder for a thorough reading of the manuscript and for helpful suggestions.

Author Contributions

G.H. devised the method, and performed the calculations. G.H. performed the analysis. G.H. and R.C.B. wrote the manuscript. All authors reviewed the manuscript.

Additional Information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Hegde, G. and Bowen, R. C. Machine-learned approximations to Density Functional Theory Hamiltonians. *Sci. Rep.* **7**, 42669; doi: 10.1038/srep42669 (2017).

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2017