

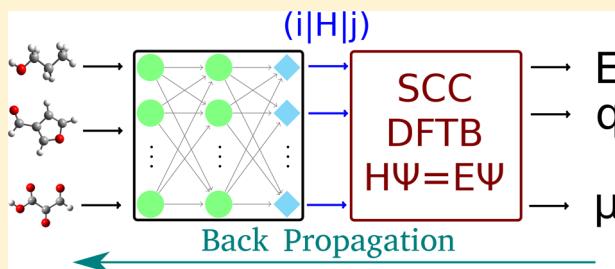
A Density Functional Tight Binding Layer for Deep Learning of Chemical Hamiltonians

Haichen Li,[†] Christopher Collins,[†] Matteus Tanha,[†] Geoffrey J. Gordon,[‡] and David J. Yaron*,[†] 

[†]Department of Chemistry, and [‡]Machine Learning Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States

 Supporting Information

ABSTRACT: Current neural networks for predictions of molecular properties use quantum chemistry only as a source of training data. This paper explores models that use quantum chemistry as an integral part of the prediction process. This is done by implementing self-consistent-charge Density-Functional-Tight-Binding (DFTB) theory as a layer for use in deep learning models. The DFTB layer takes, as input, Hamiltonian matrix elements generated from earlier layers and produces, as output, electronic properties from self-consistent field solutions of the corresponding DFTB Hamiltonian. Backpropagation enables efficient training of the model to target electronic properties. Two types of input to the DFTB layer are explored, splines and feed-forward neural networks. Because overfitting can cause models trained on smaller molecules to perform poorly on larger molecules, regularizations are applied that penalize nonmonotonic behavior and deviation of the Hamiltonian matrix elements from those of the published DFTB model used to initialize the model. The approach is evaluated on 15 700 hydrocarbons by comparing the root-mean-square error in energy and dipole moment, on test molecules with eight heavy atoms, to the error from the initial DFTB model. When trained on molecules with up to seven heavy atoms, the spline model reduces the test error in energy by 60% and in dipole moments by 42%. The neural network model performs somewhat better, with error reductions of 67% and 59%, respectively. Training on molecules with up to four heavy atoms reduces performance, with both the spline and neural net models reducing the test error in energy by about 53% and in dipole by about 25%.



1. INTRODUCTION

Machine learning (ML) has the potential to predict molecular properties at low computational cost, making it possible to rapidly search chemical space for optimal systems.^{1–5} A powerful strategy for developing ML models begins by using quantum chemistry (QC) to generate a set of data that spans the chemical space of interest.^{6–8} The ML model is then trained to reproduce the QC data with acceptable accuracy and with substantially reduced computational cost.⁹ The ML models explored to date use, as input, either features derived solely from the structure of the molecular system,^{6,9–15} or, in the case of Δ -machine-learning models,^{16–19} a combination of structural features and results from lower-cost quantum chemical methods. While quantum mechanics is used to generate the training data, the form of the ML model does not itself incorporate aspects of quantum mechanics. In the models explored here, quantum chemistry is an integral part of the prediction process. The approach may be viewed as an extension of semiempirical QC in which ML models are used to generate the empirical parameters for the model Hamiltonian.

The models explored here are developed and tested on a subset of the ANI-1 data set,²⁰ which includes small organic molecules distorted from their equilibrium positions. The ANI-1 potential²¹ is a neural network that achieves high accuracy on the entire ANI-1 data set. The network architecture is similar

to that introduced by Behler and Parrinello²² which obtains high performance on a range of chemical systems.^{23–26} Incorporating QC into the neural network may be viewed as a way to incorporate domain knowledge into the ML model. Semiempirical QC models are able to describe a broad range of chemical phenomena, including valency, bond formation, and aromaticity. Semiempirical QC methods have also had considerable quantitative success.^{27–31} Incorporating this domain knowledge into the neural network may help lower the amount of data needed to train the model and help improve transfer between systems.³² However, we note that in domains such as computer vision and natural language processing, such domain knowledge has been pruned from the ML models as the amount of training data and sophistication of the ML approaches have increased.^{33–37}

The QC portion of the model developed here is based on self-consistent-charge Density-Functional Tight-Binding (DFTB) theory.^{38–40} DFTB uses a minimal-basis, valence-electron only Hamiltonian, with atomic point charges used to describe Coulomb interactions between atoms. A unique characteristic of DFTB is that the parameters for the electronic Hamiltonian are obtained through a nonempirical approach.⁴¹ This approach

Received: August 26, 2018

Published: October 15, 2018

begins by generating QC solutions for isolated atoms placed in a simple electrostatic potential that constrains the radial spread of the electron distribution. The resulting atomic wave functions are then used to derive parameters for the electronic Hamiltonian.^{42–46} DFTB typically also includes a repulsive potential that is handled empirically by fitting to QC or experimental results for molecules or periodic systems.^{47–50}

Below, we implement DFTB as a layer for deep learning, using the TensorFlow deep learning framework.⁵¹ The DFTB layer takes, as input, values for the Hamiltonian matrix elements that define the DFTB model and generates, as output, molecular properties that are self-consistent-field solutions of the corresponding DFTB Hamiltonian. The DFTB layer supports backpropagation, allowing the model parameters to be updated efficiently. Below, training on 12 400 small organic molecules, sampled from the ANI-1 data set,²⁰ takes about 5 h on six processor cores.⁵²

Our focus here is on the construction and characterization of the DFTB layer itself, which is agnostic with respect to the form of the layers used to generate the DFTB matrix elements. To help characterize the DFTB layer, we consider two different types of earlier layers. In the spline model, we use spline functions to allow the matrix elements to be functions of the interatomic distance, r . In the neural network model, we use feed-forward neural networks (FFNN) to allow the matrix elements to be more general functions of the molecular geometry. We will refer to the combination of the input layers and the DFTB layers as a DFTB-ML model.

Because DFTB-ML models are highly flexible, we begin by initializing the input layers to the matrix elements from a published DFTB parametrization.³⁸ Training of the DFTB-ML model may be viewed as refining this initial DFTB model. Although performance on the molecules in the training set improves as training progresses, performance on test molecules may begin to degrade with continued training. Below, such overfitting is especially prevalent when the model is trained on smaller molecules and then applied to larger molecules. We consider two different regularizations to reduce such overfitting. The first applies only in the spline model and constrains the Hamiltonian matrix elements to have a monotonic dependence on r . The second regularization, which is much more effective and applies to both spline and FFNN models, penalizes deviation from the initial DFTB parameters. This penalty is a regularization that limits the model flexibility and helps in the transfer of models from smaller to larger molecules.

One motivation for considering transfer from smaller to larger molecules is that this relates to a potential advantage of building quantum mechanics directly into the ML model form. During training, we modify only matrix elements that describe short-range interactions between atoms. The longer-range interactions are described via Coulomb's law. This transition from empirical interactions at short-range to physics-based interactions at long-range may improve the degree to which models trained on small molecules are able to transfer to larger systems.

2. RELATED WORK

In current DFTB models, empirical fits are typically limited to the repulsive potential, a classical potential whose energy is added to the energy obtained from the DFTB electronic Hamiltonian. A number of approaches have been developed to help automate fits of the DFTB repulsive potential to energies and forces obtained from *ab initio* QC.^{47–50} The repulsive potential is typically written as a sum of interatomic potentials

that are nonzero over a limited range, typically just beyond the range of a covalent bond. The functional forms are also often restricted, for example, to sums of exponentials. More recently, unsupervised learning has been used to develop repulsive potentials with more general model forms.⁵³ Because our focus here is on the electronic Hamiltonian, we use spline functions to obtain a flexible, but relatively simple, form for the repulsive potential.

A number of recent DFTB parametrizations have, in addition to fitting the repulsive potential, empirically adjusted parameters that define the electronic Hamiltonian during the fitting process.^{54–58} Adjusted parameters include those that define the constraining potential and electron density cutoffs in the standard approach utilized to construct the DFTB electronic Hamiltonian from QC solutions for isolated atoms.^{54,56,57} Empirical fits have also adjusted the atomic orbital energies, and the Hubbard parameters that specify electron–electron repulsion.^{54–57} Matrix elements between atoms have also been adjusted by fitting analytic forms that describe the dependence of these matrix elements on r and that involve 2⁵⁵ or between 12 and 15⁵⁸ free parameters per matrix element type. This past work suggests that adjusting the electronic Hamiltonian can lead to significant improvements in model accuracy.

Empirical fits of the electronic Hamiltonian have used a number of optimization schemes. In a dual loop approach, optimization of the electronic Hamiltonian alternates with optimization of the repulsive potential, using different objective functions for each of these two optimization loops.^{56,57} Simultaneous optimization of all fitting parameters has also been done using gradient free optimization methods such as swarm optimization,⁵⁴ simulated annealing,⁵⁵ and pattern search.⁵⁸ These optimization methods do not require the gradients of the molecular properties with respect to model parameters to be computed. This has the advantage of allowing parameters to be adjusted to fairly complex targets such as minimum-energy structures, lattice parameters, and energy differences between polymorphs. The approach developed here uses back-propagation to efficiently compute gradients. This allows flexible models, that involve a large number of parameters, to be trained on large sets of molecular data. However, the targets are limited to molecular properties such as energy and dipole that are functions only of the input molecular geometries.

An alternative, ML-based approach, for optimizing parameters in the electronic Hamiltonian has been applied to the OM2 semiempirical QC model.^{59,60} This iterative approach considers one parameter at time. For each molecule in the data set, the value of the parameter that minimizes the error for that individual molecule is obtained. Kernel Ridge Regression (KRR) is then used to predict this optimal value, using only the molecular structure. This approach reduced mean absolute errors in atomization energies on test molecules from 6.7 to 1.3 kcal/mol.

3. METHODS

3.1. DFTB Hamiltonian. The DFTB model originates from approximations applied to the Kohn–Sham equations of density functional theory.³⁸ Here, we use the model as an empirical form for model fitting and so describe the model from an empirical perspective. The single-electron wave functions or molecular orbitals, Ψ_a , are expressed in a minimal atomic basis, ϕ_i

$$\Psi_a = \sum_i^{N_{\text{basis}}} C_{i,a} \phi_i \quad (1)$$

where the N_{basis} atomic basis functions include only valence orbitals, for example, 1s on H and 2s,2p on second-row elements.

The Kohn–Sham equations may be written as the following eigensystem,

$$\sum_j^{N_{\text{basis}}} [H_{i,j} + H_{i,j}^{(2)}(\Delta q) - \epsilon_a S_{i,j}] C_{j,a} = 0, \forall a, i \quad (2)$$

where $H_{i,j}$ and $S_{i,j}$ are matrix elements of the one-electron Hamiltonian operator and overlap operator, respectively, between atomic orbitals i and j , and ϵ_a is the Kohn–Sham orbital energies sorted from low to high. $H_{i,j}^{(2)}(\Delta q)$ describes the two-electron interactions as interactions between charge fluctuations in the atomic shells, Δq ,

$$H_{i,j}^{(2)}(\Delta q) = \frac{1}{2} S_{i,j} \sum_{\alpha}^{N_{\text{shells}}} (\gamma_{\text{shell}(i),\alpha} + \gamma_{\text{shell}(j),\alpha}) \Delta q_{\alpha} \quad (3)$$

where α indexes over the N_{shells} atomic shells (e.g., 1s,2s,2p) in the molecule, $\text{shell}(i)$ is the shell of atomic orbital i , Δq_{α} is the charge fluctuation of shell α , and $\gamma_{\beta,\alpha}$ is the Coulomb interaction between shells β and α .

The charge fluctuations are obtained from the Mulliken population of each atomic shell. These may be obtained from the electronic density matrix,

$$\rho_{i,j} = \sum_a n_a C_{i,a} C_{j,a} \quad (4)$$

where i and j index atomic orbitals and n_a is the occupation of the a th molecular orbital. The Mulliken charge of each atomic shell is then

$$q_{\alpha} = - \sum_{i \in \text{shell}(\alpha)} \sum_j \rho_{i,j} S_{i,j} \quad (5)$$

The charge fluctuation of the α th shell is then,

$$\Delta q_{\alpha} = q_{\alpha} - q_{\alpha}^{(0)} \quad (6)$$

where $q_{\alpha}^{(0)}$ is the charge of atomic shell α in the isolated, neutral, atom. The values of $q_{\alpha}^{(0)}$ are constants that may be taken as part of the model parametrization.

In section 3.3 below, we convert $H^{(2)}(\Delta q)$ from the summation over atomic shells in eq 3 to a summation over atomic orbitals. With each atomic orbital ϕ_i , we introduce a charge

$$Q_i = - \sum_j \rho_{i,j} S_{i,j} \quad (7)$$

and a charge fluctuation

$$\Delta Q_i = Q_i - Q_i^{(0)} \quad (8)$$

where the reference charge for the atomic orbitals, $Q^{(0)}$, is obtained by distributing the reference charge for the shells, $q^{(0)}$ of eq 6, equally across all atomic orbitals in the respective shell. $H_{i,j}^{(2)}(\Delta q)$ of eq 3 may then be written

$$H_{i,j}^{(2)}(\Delta q) = \frac{1}{2} S_{i,j} \sum_k (G_{i,k} + G_{j,k}) \Delta Q_k \quad (9)$$

with

$$G_{i,j} = \gamma_{\text{shell}(i),\text{shell}(j)} \quad (10)$$

The total electronic energy is then

$$E_{\text{elec}} = \sum_{i,j} \rho_{i,j} H_{i,j} + \frac{1}{2} \sum_{i,j} \Delta Q_i G_{i,j} \Delta Q_j \quad (11)$$

The total energy of the system also includes a classical potential energy term referred to as the repulsive potential because it is intended to include the repulsive interaction between the atomic cores that are not included in the electronic Hamiltonian

$$E_{\text{rep}} = \sum_{A>B} R_{Z_A, Z_B} (|\mathbf{r}_A - \mathbf{r}_B|) \quad (12)$$

where A and B label atoms, \mathbf{r}_A is the Cartesian position of the A th atom, and R_{Z_A, Z_B} is a function that depends on the elements, indicated by atomic numbers Z_A , of the atoms.

We also include a reference energy such that the total energy is given by,

$$E_{\text{tot}} = E_{\text{elec}} + E_{\text{rep}} + E_{\text{ref}} \quad (13)$$

DFTB parametrizations typically consider only the change in energy due to geometric distortion.⁵⁰ This can be done, for example, by fitting to the energy differences between molecular geometries or by fitting to atomic forces. Fitting to distortion energies may be viewed as assigning a separate reference energy to each isomer under consideration. Below, we instead use a reference energy based on atom counts,

$$E_{\text{ref}} = \sum_{Z=1,6,8} p_Z N_Z + p_c \quad (14)$$

where the sum is over the elements (H,C,O), p_Z is a parameter associated with each element, N_Z is the number of occurrences of that element in the molecule, and p_c is a single parameter that sets the overall zero of energy. Equation 14 has the advantage of allowing the model to predict energy differences between any molecules composed of elements present in the training data.

3.2. Matrix Element Models. The parameters of a DFTB model are used to construct the matrices \mathbf{H} , \mathbf{S} , and \mathbf{G} of eqs 2 and 9, and the repulsive potentials, R_{Z_A, Z_B} of eq 12. In practice, a DFTB parametrization is specified through a set of files that list constants for the on-atom terms and provide the between-atom terms as either a parametric or tabulated function of the interatomic separation, r . Our goal is to provide an efficient means to learn more flexible forms for these parameters. Below, we use either splines or feed-forward neural networks as examples of such flexible forms. However, our intent is to support any model form that can be implemented efficiently in a deep-learning framework such as Tensorflow. We therefore divide the responsibilities between “models” that generate information on the type currently extracted from the files that define a DFTB parametrization, and a “DFTB layer” that uses this information to predict molecular properties. The models are implemented as layers that provide input to the DFTB layer. The DFTB layer is agnostic with regards to the form of these previous layers, requiring only that these earlier layers produce the information in the order specified in section 3.3.

Although our implementation allows deep learning to be used to construct \mathbf{S} , we currently fix \mathbf{S} to values from an existing DFTB parametrization³⁸ and so do not explore the potential benefits and challenges associated with empirically fitting \mathbf{S} . Also, because our focus is on the electronic Hamiltonian,

we restrict the repulsive potentials, $R_{Z_A Z_B}$ of eq 12, to one-dimensional spline functions.

Table 1 lists the models needed to construct the matrices \mathbf{H} and \mathbf{G} for molecules composed of the elements H, C and O.

Table 1. Types of Aligned DFTB Matrix Elements^a

operator	block type	elements	shell type	models for C,H,O
\mathbf{H}	diagonal	Z	s, p	5
\mathbf{H}	off-diagonal	Z_1, Z_2	ss, sp, pp_{σ} , pp_{π}	18
\mathbf{G}	diagonal	Z	ss, pp, sp	7
\mathbf{G}	off-diagonal	Z_1, Z_2	ss, pp, sp	15

^aBlock type refers to blocks of the operator on a single atom (diagonal) or between two atoms (off-diagonal). Shell type refers to the orbitals involved in the aligned matrix elements. Blocks of the operators on or between atoms can be generated from these aligned matrix elements through Slater–Koster rotations. The last column lists the number of models of each type needed for molecules containing the elements C, H, and O.

These matrices are assembled from “blocks”, with diagonal blocks being between orbitals on a single atom and off-diagonal blocks being between orbitals on two different atoms. For diagonal blocks of \mathbf{H} , the models generate the energy of the s and p atomic orbitals. We allow these orbital energies to depend on the environment of the atom,⁶¹ but assume that all orbitals of a given shell on a given atom have the same energy. For off-diagonal blocks of \mathbf{H} , the models generate the matrix elements between atomic orbitals that are aligned along the axis connecting the two atoms. We will refer to these as “aligned” matrix elements. For first-row elements, the unique aligned elements are ss, sp, pp_{σ} and pp_{π} . These aligned matrix elements may depend on the environments of the associated atoms. Slater–Koster (SK) rotations are linear transformations that rotate these aligned elements into blocks of \mathbf{H} between atoms. The DFTB layer described below carries out the SK rotations, with the input layers generating only the aligned matrix elements. For both diagonal and off-diagonal blocks, the models for \mathbf{G} generate the Coulomb interaction, γ of eq 3, between the various shells, ss, sp, and pp. Because Coulomb interactions between atoms in DFTB are between point charges associated with each atomic shell, eq 3, SK rotations are not necessary for \mathbf{G} .

3.3. Expressing DFTB as Tensor Operations. Training of the DFTB model requires a large number of gradient descent updates to the model parameters. To make these updates efficient, we implement DFTB as a series of tensor operations in the Tensorflow deep learning framework.⁵¹ Each update to the parameters is based on a set of molecules that we will refer to as a minibatch. In Tensorflow, a computational graph is first constructed that specifies the series of tensor operations required to make predictions for a single minibatch. Data that depends on the particular molecules in the minibatch are then fed into this graph. In a forward pass through the graph, the molecular properties are predicted based on the current model parameters, and the difference between predicted and target properties is used to compute a loss. In a backward pass, the gradient of the loss with respect to model parameters is computed and used to update the model parameters. Here, we use Tensorflow’s ADAM optimizer⁶² to update the parameters based on the computed gradients, with a learning rate of 10^{-5} , first moment exponential decay rate of $\beta_1 = 0.9$, second moment exponential decay rate of $\beta_2 = 0.999$, and numerical stability constant $\epsilon = 10^{-8}$.

An epoch of training corresponds to a forward and backward pass performed on each minibatch in the training data. The form of the computational graph depends only on the empirical formulas of the molecules included in the minibatch. To allow us to use a single graph during model training, the molecules in each minibatch have identical empirical formulas, and the atoms in each molecule are sorted to have identical orders of elements. (This restriction on the empirical formulas of the minibatches could be relaxed by sharing model parameters between multiple computational graphs.)

As discussed in section 3.2, the input to the DFTB layer comes from models that produce the aligned matrix elements of Table 1 in a specific order. This order is specified during construction of the computational graph and depends only on the sequence of empirical formulas in a minibatch. The specification consists of a list of tuples that specify the type of aligned matrix element (as in Table 1), the index of the molecule within the minibatch, and the indices of the respective atoms within that molecule. The layers that provide input to the DFTB layer may use these specifications to implement a wide variety of models. The list is ordered by type of aligned matrix element so that the input layers may produce all matrix elements of a given type with a single set of tensor operations. Consider, for example, the use of a separate feed-forward neural network for each of the matrix elements types, or models, in Table 1. From the molecular geometries in a particular minibatch, a list of feature vectors may be created that, when fed into the neural network for a model, produces the aligned matrix elements in the specified order. For diagonal elements, the feature vectors may describe the environment of the atom associated with that matrix element. For off-diagonal elements, the feature vectors may describe the environment of the pair of atoms associated with that matrix element.

The initial operations in the DFTB layer carry out the Slater–Koster (SK) rotations (Figure 1b). The rotations are

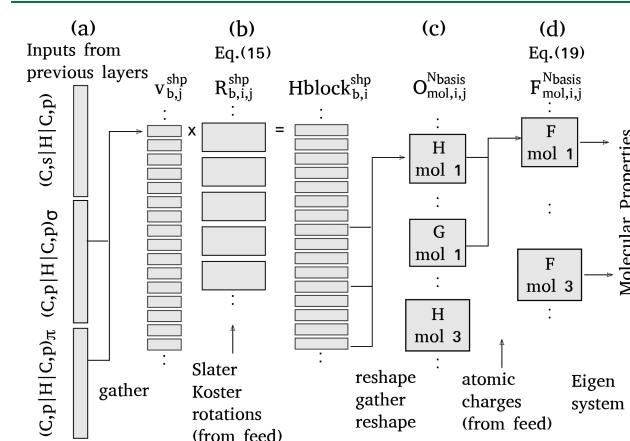


Figure 1. Schematic representation of the DFTB layer. (a) The aligned matrix elements are input from previous layers, ordered as specified during construction of the computational graph. (b) Tensorflow gather and reshape operations are used to rearrange the input matrix elements and to carry out Slater–Koster rotations. (c) The resulting operator blocks are assembled into operators for each molecule. (d) Fock operators are constructed using atomic charges fed into the graph, and molecular properties are predicted and output from the layer.

written as batch matrix multiples, which can be handled efficiently in Tensorflow or other deep-learning framework,

$$\mathbf{Hblock}_{b,i}^{\text{shp}} = \sum_j \mathbf{R}_{b,i,j}^{\text{shp}} \mathbf{v}_{b,j}^{\text{shp}} \quad (15)$$

The superscript shp refers to the shape of the individual SK rotation matrices, that is, the dimensions of indices i and j in eq 15. For example, shp = 3×1 arises when rotating the single aligned element of orbital type sp to the three matrix elements in the operator block between s and p orbitals. shp = 9×2 arises when rotating the two aligned elements pp $_{\sigma}$ and pp $_{\pi}$ to the 3×3 operator block between p orbitals. The index b in eq 15 labels blocks of the H operator. For computational efficiency, a single evaluation of eq 15 carries out all required SK rotations of a given shape, that is, b in eq 15 runs over all blocks in the minibatch requiring SK rotations with shape shp. This requires substantial rearrangements of the values input to the DFTB layer. These are done using Tensorflow's *gather* and *reshape* operators. Given a vector \mathbf{x} and a list of integers \mathbf{I} , *gather* returns $y_j = x_{I_j}$. *Reshape* transforms a N -dimensional tensor into a flattened 1-dimensional form and vice versa. For each shape, shp in eq 15, a flattened view of $v_{b,j}^{\text{shp}}$ is gathered from the input and reshaped into a 2-dimensional tensor. The batch matrix multiply of eq 15 then generates flattened views of the operator blocks, $\mathbf{Hblock}_{b,i}^{\text{shp}}$, ordered by a block index b .

An analogous approach is used to assemble blocks of \mathbf{G} . However, because SK rotations are not needed for \mathbf{G} , the values are only gathered and reshaped into $\mathbf{Gblock}_{b,i}^{\text{shp}}$.

The next operations in the DFTB layer assemble the operator blocks into matrices for the operators \mathbf{H} and \mathbf{G} of each molecule in the minibatch (Figure 1c). For computational efficiency, all operators of a given dimension, N_{basis} , are assembled through a single gather operation. The operator blocks, $\mathbf{Hblock}_{b,i}^{\text{shp}}$ of eq 15, are first flattened and concatenated into a single vector that holds results for all shapes, shp. For each N_{basis} , a flattened version of all operators with dimension N_{basis} are then gathered and reshaped into $\mathbf{O}_{\text{mol},i,j}^{N_{\text{basis}}}$, where mol indexes molecules and \mathbf{O} includes both \mathbf{H} and \mathbf{G} .

At this point, we have transformed the aligned matrix elements input to the DFTB layer into the operator matrices needed for the Kohn–Sham equations of eq 2.

We next compute the two-electron contributions to the Kohn–Sham equations, $\mathbf{H}^{(2)}$ of eq 2, with a tensor operation for each value of N_{basis} ,

$$\mathbf{H}_{\text{mol},i,j}^{(2)N_{\text{basis}}} = \frac{1}{2} \mathbf{S}_{\text{mol},i,j}^{N_{\text{basis}}} \sum_k (\mathbf{G}_{\text{mol},i,k}^{N_{\text{basis}}} + \mathbf{G}_{\text{mol},j,k}^{N_{\text{basis}}}) \Delta \mathbf{Q}_{\text{mol},k}^{N_{\text{basis}}} \quad (16)$$

The charge fluctuations, $\Delta \mathbf{Q}_{\text{mol},k}$, are initialized from the starting model parameters and updated to obtain self-consistency, as discussed below.

The generalized eigenvalue problem of eq 2 is converted to a self-adjoint eigenvalue problem. The overlap matrices are first diagonalized

$$\mathbf{S}_{\text{mol},i,j}^{N_{\text{basis}}} = \sum_k \mathbf{U}_{\text{mol},i,k}^{N_{\text{basis}}} \Lambda_{\text{mol},k}^{N_{\text{basis}}} \mathbf{U}_{\text{mol},j,k}^{N_{\text{basis}}} \quad (17)$$

and the results are used to form,

$$\Phi_{\text{mol},i,j}^{N_{\text{basis}}} = \mathbf{U}_{\text{mol},i,j}^{N_{\text{basis}}} (\Lambda_{\text{mol},j}^{N_{\text{basis}}})^{-1} \quad (18)$$

A Fock operator for the self-adjoint eigensystem is then formed,

$$\mathbf{F}_{\text{mol},i,j}^{N_{\text{basis}}} = \sum_{k,l} \Phi_{\text{mol},k,i}^{N_{\text{basis}}} (\mathbf{H}_{\text{mol},k,l}^{N_{\text{basis}}} + \mathbf{H}_{\text{mol},l,k}^{(2)N_{\text{basis}}}) \Phi_{\text{mol},l,j}^{N_{\text{basis}}} \quad (19)$$

and diagonalized

$$\mathbf{F}_{\text{mol},i,j}^{N_{\text{basis}}} = \sum_k \mathbf{C}_{\text{mol},i,k}^{N_{\text{basis}}} \epsilon_{\text{mol},k}^{N_{\text{basis}}} \mathbf{C}_{\text{mol},j,k}^{N_{\text{basis}}} \quad (20)$$

The tensors $\epsilon_{\text{mol},k}^{N_{\text{basis}}}$ hold the orbital energies, ϵ of eq 2, and the orbital expansion coefficients, \mathbf{C} of eq 1, are given by

$$\mathbf{C}_{\text{mol},i,j}^{N_{\text{basis}}} = \sum_k \Phi_{\text{mol},i,k}^{N_{\text{basis}}} \mathbf{C}_{\text{mol},j,k}^{N_{\text{basis}}} \quad (21)$$

The density matrix is obtained by first masking out the unoccupied orbitals,

$$\mathbf{Cocc}_{\text{mol},i,j}^{N_{\text{basis}}} = \mathbf{Mask}_{\text{mol},i,j}^{N_{\text{basis}}} \mathbf{C}_{\text{mol},i,j}^{N_{\text{basis}}} \quad (22)$$

where $\mathbf{Mask}_{\text{mol},i,j}^{N_{\text{basis}}}$ is 1 for occupied orbitals j , and 0 otherwise. The density matrices are then

$$\rho_{\text{mol},i,j}^{N_{\text{basis}}} = 2 \sum_k \mathbf{Cocc}_{\text{mol},i,k}^{N_{\text{basis}}} \mathbf{Cocc}_{\text{mol},j,k}^{N_{\text{basis}}} \quad (23)$$

The mask of eq 22 allows tensor operations to be used despite molecules with the same N_{basis} having potentially different numbers of occupied molecular orbitals. The charge fluctuations, per atomic orbital, are then obtained from

$$\Delta \mathbf{Q}_{\text{mol},i}^{N_{\text{basis}}} = - \sum_k \mathbf{S}_{\text{mol},i,k}^{N_{\text{basis}}} \rho_{\text{mol},i,k}^{N_{\text{basis}}} - \mathbf{Q}_{\text{mol},i}^{(0)N_{\text{basis}}} \quad (24)$$

where $\mathbf{Q}_{\text{mol},i}^{(0)N_{\text{basis}}}$ holds the reference charges, $\mathbf{Q}^{(0)}$ of eq 8. The electronic energy is then,

$$\mathbf{E}_{\text{elecmol}}^{N_{\text{basis}}} = \sum_{i,j} \rho_{\text{mol},i,j}^{N_{\text{basis}}} \mathbf{H}_{\text{mol},i,j}^{N_{\text{basis}}} \quad (25)$$

$$+ \frac{1}{2} \sum_{i,j} \Delta \mathbf{Q}_{\text{mol},i}^{N_{\text{basis}}} \mathbf{G}_{\text{mol},i,j}^{N_{\text{basis}}} \Delta \mathbf{Q}_{\text{mol},j}^{N_{\text{basis}}} \quad (26)$$

The molecular dipoles are obtained from $\Delta \mathbf{Q}^{N_{\text{basis}}}$ and the Cartesian positions of the atoms.

Because the overlap matrices are not altered during model training, $\Phi_{\text{mol},i}^{N_{\text{basis}}}$ may be computed once at the start of training and used throughout. However, because the overlap matrices depend on molecular geometry, $\Phi_{\text{mol},i}^{N_{\text{basis}}}$ must be computed separately for each minibatch.

Given the computational graph for a forward pass through the DFTB layer, Tensorflow computes the gradients needed to train the model parameters. If the molecular orbitals become degenerate, the gradients associated with the eigensystem diverge. Any values in the gradient evaluation that are returned as undefined are set to zero. In assigning molecular geometries of a given isomer to train and test sets, we also sort such that train molecules are less likely to encounter degeneracies (see Supporting Information). Because of this, undefined values for the gradients occur very rarely.

3.4. Self-Consistent Field. The two-electron portion of the Hamiltonian depends on the current estimates for the atomic charge fluctuations, $\Delta \mathbf{Q}$ of eq 16. In application of DFTB to a new molecule, these charges are obtained through an iterative procedure that locates a fixed point, in which the charges used to construct $\mathbf{H}^{(2)}$ in eq 16 agree with those generated from the use of $\mathbf{H}^{(2)}$, that is, those predicted by eq 24. In the DFTB layer of section 3.3, the charge distributions used to construct $\mathbf{H}^{(2)}$ are fed into the computational graph. Updating of the charges is therefore handled outside of the Tensorflow computational graph.

Roothaan–Hall is a simple iterative approach in which each iteration uses the charges generated by the previous iteration to update $\mathbf{H}^{(2)}$. Although such an approach is simple to implement, our experience is that this approach often fails to converge. We instead use the direct inverse of the iterative subspace (DIIS) method.⁶³ Periodically during training, the current \mathbf{H} and \mathbf{G} operator matrices are exported to a module that uses DIIS iterations to obtain the SCF charge distribution. The resulting charges are then fed into the computational graph, for use in constructing $\mathbf{H}^{(2)}$. Figure 2 shows a representative

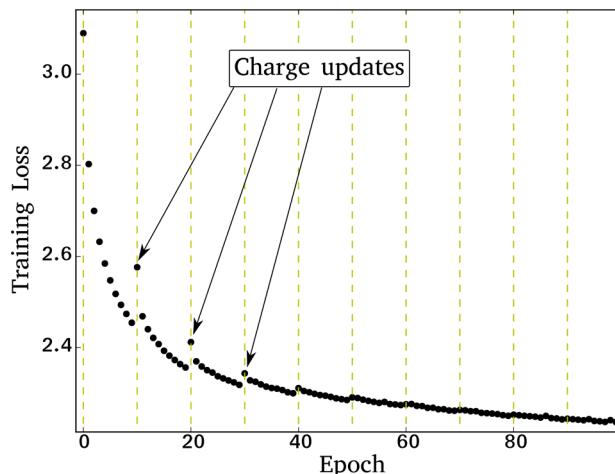


Figure 2. Training error as a function of number of epochs. Every 10 epochs, the SCF solution for the current Hamiltonian parameters is used to update the charges fed into the DFTB layer. (Results are for a spline model with no regularization trained on molecules with up to seven heavy atoms.).

training example in which the charges are updated every 10 epochs. Between charge updates, training brings the predictions closer into agreement with the target values, but this training ignores the impact that changes to the parameters have on the charges used to construct $\mathbf{H}^{(2)}$. When the charges are updated, the agreement between predicted and target values degrades as shown by the spike in the loss at 10 and 20 epochs in Figure 2. As training continues and the model parameters begin to stabilize, updates to charges have smaller impacts on the loss. For the remainder of this paper, we

perform charge updates every 10 epochs and show only results obtained immediately following a charge update. These are the results that would be obtained if the current model parameters were used to obtain SCF solutions and so are an accurate reflection of model performance. Charge updates account for roughly half of the total computation time.

In our initial exploration of model forms, DIIS occasionally failed to converge. In such cases, we did not update the feed to the DFTB computational graph. The charges fed into the computational graph for such molecules were then those from the most recent converged DIIS procedure. While this approach allowed us to explore a wide variety of model forms, for the results presented here, DIIS converged in all cases.

4. DATA SET

The results presented here are for molecules composed of the elements H, C, and O from the ANI-1 data set.²¹ The molecular structures in this data set were obtained by the Normal Mode Sampling (NMS) method²¹ and reflect the types of structures that may arise in a room temperature simulation. Because the ANI-1 data set does not contain atomic charges, we used the GAUSSIAN program⁶⁴ to generate data for each of the included molecules using the same quantum chemical method as used in ANI-1, that is, Density Functional Theory with the ω B97X functional and 6-31g(d) basis set. Electrostatic potential, ESP, charges were obtained using the technique of Hu et al., which gives charges that vary smoothly with changes in molecular geometry and basis set.⁶⁵

As discussed in section 3.3, the computational graph depends on the empirical formulas of the molecules in each minibatch. We use a single computational graph corresponding to the 66 empirical formulas of Table 2. The formulas are sorted by number of heavy atoms, with 10 molecules each for heavy atom counts of 1 to 5. For larger systems, the number of molecules in the minibatch drops roughly as the square of the number of heavy atoms. Because the number of aligned matrix elements in a molecule scales roughly quadratically with the number of heavy atoms, this choice distributes the aligned matrix elements computed for molecules with five or more heavy atoms roughly equally among different size molecules. The Supporting Information provides additional details on the selection of molecular geometries to include in the data set.

Table 2. Data Used to Train the Model Consists of 200 Train and 50 Test Minibatches^a

no. of heavy atoms	empirical formulas	no. of molecules		
		minibatch	train set	test set
1	H_2O CH_4 H_2O CH_4 H_2O CH_4 H_2O CH_4 H_2O CH_4	10	2000	500
2	H_2O_2 CH_4O CH_2O C_2H_2 C_2H_6 C_2H_4 H_2O_2 CH_4O CH_2O C_2H_2	10	2000	500
3	$\text{C}_2\text{H}_4\text{O}$ $\text{C}_2\text{H}_6\text{O}$ C_3H_6 H_2O_3 CO_2 CH_2O_2 C_3H_4 C_3H_8 $\text{C}_2\text{H}_4\text{O}$ $\text{C}_2\text{H}_6\text{O}$	10	2000	500
4	$\text{C}_3\text{H}_6\text{O}$ C_4H_8 $\text{C}_2\text{H}_4\text{O}_2$ $\text{C}_3\text{H}_8\text{O}$ $\text{C}_3\text{H}_4\text{O}$ C_4H_6 CH_2O_3 $\text{C}_2\text{H}_2\text{O}_2$ $\text{C}_2\text{H}_6\text{O}_2$ $\text{C}_3\text{H}_2\text{O}$	10	2000	500
5	$\text{C}_4\text{H}_8\text{O}$ $\text{C}_4\text{H}_6\text{O}$ $\text{C}_3\text{H}_6\text{O}_2$ C_3H_{10} $\text{C}_3\text{H}_4\text{O}_2$ C_5H_8 $\text{C}_4\text{H}_{10}\text{O}$ C_5H_6 $\text{C}_3\text{H}_8\text{O}_2$ $\text{C}_4\text{H}_4\text{O}$	10	2000	350
6	$\text{C}_5\text{H}_{10}\text{O}$ $\text{C}_5\text{H}_8\text{O}$ $\text{C}_4\text{H}_8\text{O}_2$ $\text{C}_4\text{H}_6\text{O}_2$ $\text{C}_5\text{H}_6\text{O}$ C_6H_{10} C_6H_8	7	1400	350
7	$\text{C}_6\text{H}_{10}\text{O}^*$ $\text{C}_5\text{H}_8\text{O}_2^*$ $\text{C}_6\text{H}_8\text{O}$ $\text{C}_6\text{H}_2\text{O}$ $\text{C}_5\text{H}_{10}\text{O}_2$	5	1000	250
8	$\text{C}_6\text{H}_{10}\text{O}_2^*$ $\text{C}_7\text{H}_{10}\text{O}^*$ $\text{C}_7\text{H}_1\text{I}_2\text{O}^*$ $\text{C}_6\text{H}_8\text{O}_2^*$	4		200
Total:		66	12400	3300

^aEach minibatch has 66 molecules with the indicated empirical formulas. For empirical formulas with greater than 250 isomers, indicated with an asterisk (*), the isomers in the test set are distinct from those in the training set.

5. LOSS BEING MINIMIZED

The model parameters are updated to minimize a weighted sum of the root-mean-square (RMS) error in the target molecular properties,

$$\text{Loss} = \sum_{\text{prop}} w_{\text{prop}} \sqrt{\frac{1}{N_{\text{prop}}} \sum_i^N |\text{Pred}_i - \text{Target}_i|^2} \quad (27)$$

where w_{prop} is the weight for property prop , N_{prop} is the number of predicted quantities, Pred_i is the value predicted from the current DFTB parameters and Target_i is the target value from the DFT computations of section 4.

The optimization begins with the DFTB parameters of ref 38, which we will refer to as mio-0-1. The mio-0-1 parameter set is the first freely available DFTB parameter set developed for organic molecules including O, N, C, H.^{38,66} The initial errors are shown in Table 3. The error in energy depends on how the

Table 3. RMS Errors Using Initial DFTB Parameters^a

no. of heavy atoms	E_{mol}		E_{atom}		atomic charges	actual dipole	ESP dipole
	up to 4	up to 7	up to 4	up to 7			
1	3.51	4.16	1.95	2.30	0.20	0.228	0.244
2	4.90	4.97	2.63	2.56	0.09	0.208	0.216
3	6.99	7.12	2.57	2.62	0.14	0.154	0.152
4	5.87	5.64	1.82	1.69	0.13	0.257	0.255
5	7.86	7.38	1.99	1.67	0.15	0.233	0.233
6	9.33	8.50	2.03	1.62	0.15	0.252	0.254
7	12.22	10.79	2.18	1.71	0.17	0.313	0.312
8	14.73	12.94	2.27	1.80	0.17	0.341	0.340

^a E_{mol} and E_{atom} are from fitting the reference energy, E_{ref} of eq 13, to total energy versus energy per heavy atom. Columns “up to 4” and “up to 7” are from fitting E_{ref} to molecules with up to four versus up seven heavy atoms. Dipole errors are for Cartesian components of the dipole, with *ESP* indicating that the dipole is computed from the atomic charges. Units are kcal/mol for energy, e^- for charges, and Debye for dipole.

parameters for the reference energy of eq 13 are obtained. In evaluating models below, we will fit either to molecules with up to four heavy atoms, or up to seven heavy atoms. The initial errors for these are labeled “up to 4” and “up to 7” in Table 3. In addition, the reference energy can be fit to either the total molecular energy, E_{mol} , or the energy per heavy atom, E_{atom} . The initial error in E_{mol} increases substantially with molecular size while that in E_{atom} is less dependent on size. In the fits shown below, we therefore use E_{atom} in the loss function. It is also worth noting that, even when E_{ref} is the only aspect of the model that is trained, performance for the energy is substantially enhanced by including molecules with up to seven heavy atoms.

The target properties include the ESP charges and the Cartesian components of the dipole moment computed from these ESP charges. Fitting to the ESP dipoles prevents competition between the dipole loss and the atomic charge loss. The DFTB Hamiltonian is also based on point charges and using the ESP dipole prevents the fitting process from considering aspects of the dipole moment that cannot be accounted for within DFTB. Use of the ESP dipole has only minor effects, as indicated by the close agreement between initial error of the ESP dipole and actual dipole in Table 3.

For all fits reported here, w_{prop} is $1/(0.1 \text{ kcal/mol})$ for E_{atom} , $1/(0.01)$ for charges, and $1/(0.01 \text{ D})$ for dipole components. These values were chosen based on the relative magnitudes of the initial errors in Table 3.

6. SPLINE MODEL

The DFTB layer of section 3.3 may be used to develop an semiempirical Hamiltonian in which the matrix elements depend only on the atomic elements and the distance between the atoms. The resulting model is then directly compatible with current implements of DFTB.⁶⁷ In this approach, the diagonal matrix elements of Table 1 depend only on the atomic element and orbital type while the off-diagonal elements add a dependence on interatomic distance, r . The dependence on r is handled through two cubic splines that meet at r_c , a cutoff distance that varies with element types. During training of the model, only the region below r_c is varied. Each of the two joined splines has 20 knots, with parameters initialized by a least-squares fit to the parameters of mio-0-1.³⁸ The boundary conditions at r_c force a continuous zeroth and first derivative, with natural boundary conditions used at the outside extremes.

The values for r_c were selected based on the distribution of interatomic distances in the data set of section 4 (Figure 3).

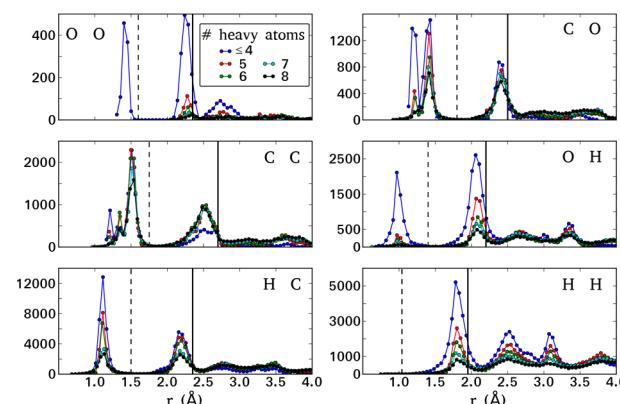


Figure 3. Distribution of interatomic distances in the training and test data of section 4, for different pairs of elements. x-Axis is interatomic distance in angstroms and y-axis is number of occurrences in interval centered on that point. Vertical lines are cutoff distances for modifications to the repulsive potential, R , (dashed) and the electronic matrix elements, H and G (solid).

For the repulsive potential, R , we use cutoffs just beyond the first peak in the distributions, such that the repulsive potential is varied only within distances corresponding to a covalent bond. This is consistent with the use of R to model repulsive interactions between atomic cores. For the electronic matrix elements, H and G , we choose r_c just beyond the second peak in the distributions of Figure 3. This allows the training of the model to modify next-nearest-neighbor electronic interactions between atoms. The model parameters are initialized to the matrix elements of mio-0-1.^{38,41,68} For R , the mio-0-1 matrix elements drop to near zero at r_c and so the boundary conditions at r_c constrain both the repulsive potential and its derivative to zero at r_c . For H and G , the constraints at r_c ensures a smooth transition to mio-0-1 values. For G , the mio-0-1 matrix elements tend toward Coulomb's law at long distances. The use of a constrained spline for G , initialized to mio-0-1 values beyond r_c , ensures that charge-charge interactions follow Coulomb's law at long distances. Transitioning

from an empirical to physics-based model for charge–charge interactions may allow models trained on smaller molecules to transfer to larger molecules. The lower and upper ranges of the splines are set by the range of interatomic separations in Figure 3.

The splines are implemented using a B-spline basis.⁶⁹ Given the values at which the spline is to be evaluated, r_i , the values, V_i , can be obtained from a single matrix multiply,

$$\mathbf{V} = \mathbf{B}\mathbf{X} + \mathbf{V}_0 \quad (28)$$

where \mathbf{X} is a vector holding the parameters that may be varied to span the space of all cubic spline functions consistent with the above boundary conditions. \mathbf{V}_0 is a constant vector that is needed because the desired boundary conditions lead to a general linear dependence on model parameters.

When the model is trained on molecules with up to seven heavy atoms, the performance on test molecules with between one and eight heavy atoms is improved, relative to the mio-0-1 values (lower panels of Figure 4). The RMS errors in energies

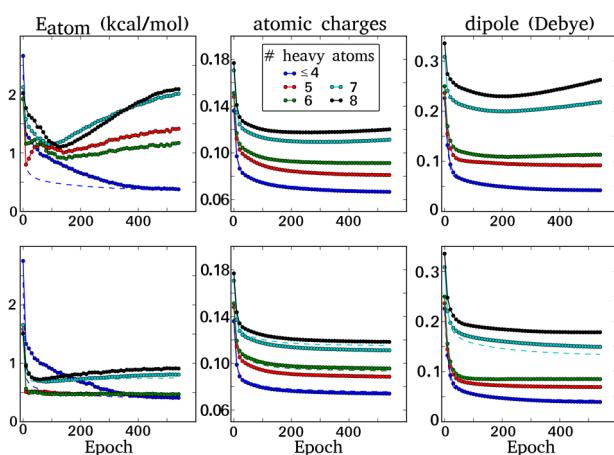


Figure 4. Spline model without regularization. Upper and lower panels are trained on molecules with up to four and seven heavy atoms, respectively. Solid lines are on test data and dotted lines are for train data. Charges are updated using DIIS to solve the self-consistent-field problem every 10 epochs, and data shown is following such updates.

and dipole components are reduced by about a factor of 2. However, when trained on molecules with up to four heavy atoms, the performance on molecules with more than four heavy atoms improves in early epochs but then begins to degrade (upper panels of Figure 4). Figure 5 shows how the spline evolves with training epoch, for matrix elements of H between p orbitals of carbon atoms aligned in a σ orientation. Similar results are obtained for other matrix elements (see Supporting Information). When trained on molecules with up to four heavy atoms, the spline begins to oscillate (Figure 5a) as the model begins to overtrain. When trained on molecules with up to seven heavy atoms, the spline has a smoother dependence on r (Figure 5d). This suggests that overfitting may be due to overly complex forms for the dependence of the matrix elements on interatomic distance, r .

To help prevent overfitting, we tried imposing a penalty on oscillatory behaviors. This was done by adding an additional term to the loss of eq 27 that penalizes nonmonotonic behavior,

$$\text{Loss}_{\text{mono}} = \lambda_{\text{mono}} \sum_i \max(0, pV'_i)^2 \quad (29)$$

where λ_{mono} sets the magnitude of the penalty, the sum is over a uniform grid with the same range as the splines defining the

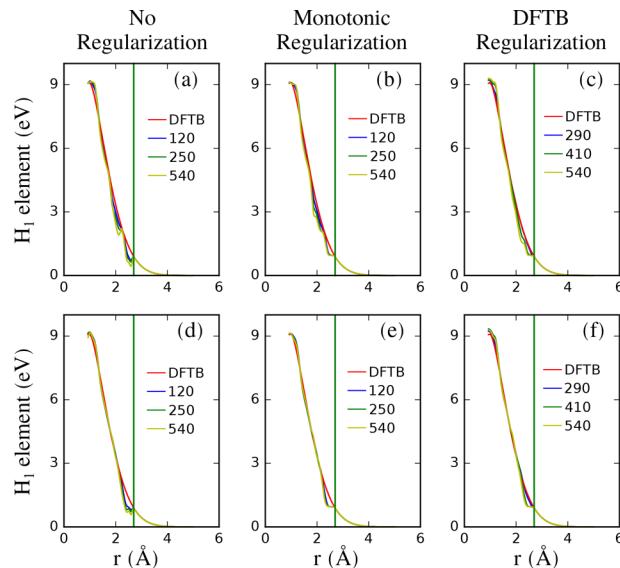


Figure 5. Dependence on interatomic separation, r , of the one-electron Hamiltonian matrix element between p orbitals of Carbon atoms, aligned in a σ overlap orientation. Upper/lower panels are for models trained with up to 4 and up to 7 heavy atoms, respectively. The right-most column is with no regularization, the middle column is with monotonic regularization, and the left-most column is for DFTB regularization.

matrix elements but with three times the density, V'_i is the derivative of the spline at the i th point of this dense grid, and p is $+1$ if the matrix element should be a decreasing function of r and -1 otherwise. This is implemented efficiently by using a linear form similar to that of eq 28 to evaluate V'_i and Tensorflow's *ReLU* function for the *max* term. Values for λ_{mono} of 10^3 and 10^5 give equivalent results, indicating that for both values, the penalty is sufficient to enforce monotonic behavior. Results shown in this paper are with $\lambda_{\text{mono}} = 10^5$. When only the repulsive potential R is allowed to vary, monotonic regularization is sufficient to lead to improved transfer from smaller to larger molecules (middle panels of Figure 6). However, when the electronic Hamiltonian is also allowed to vary, monotonic regularization is not sufficient

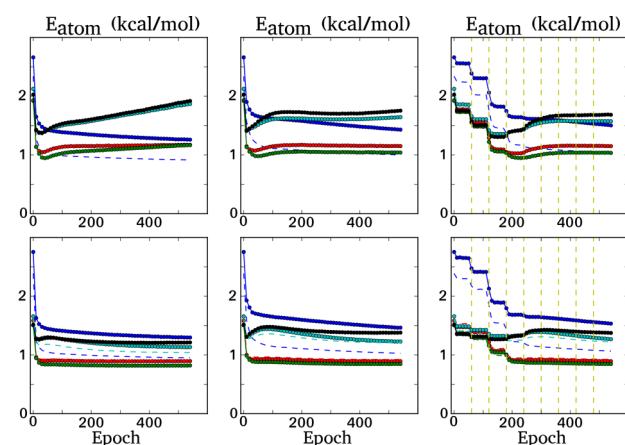


Figure 6. Spline model used only for the repulsive energy. Upper and lower panels are for training on up to 4 and up to 7 heavy atoms, respectively. Left column uses no regularization and middle column uses monotonic regularization. Right column uses DFTB regularization, with dotted lines indicating points at which the penalty for deviation from mio-0-1 starting values is relaxed.

to improve transfer to larger molecules (Figure 7). Figure 5b suggest that, even with monotonic regularization, the dependence on r remains overly complex.

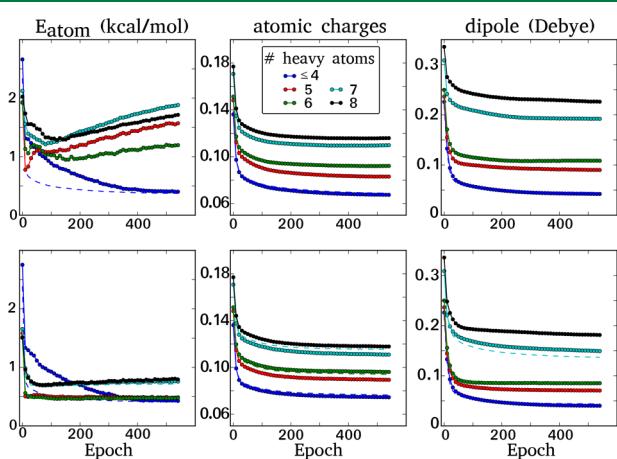


Figure 7. Spline model with regularization used to force monotonic decay, using $\lambda_{\text{mono}} = 10^5$ in eq 29. Conventions are as in Figure 4.

As a means to further constrain the model space, we added an additional penalty for deviation from the original DFTB model.

$$\text{Loss}_{\text{DFTB}} = \frac{1}{\lambda_{\text{DFTB}}^2} \frac{1}{N_{\text{elements}}} \sum_i^{N_{\text{elements}}} (V_i - \text{DFTB}_i)^2 \quad (30)$$

where λ_{DFTB} sets the magnitude of the penalty, the sum is over all matrix elements required for the minibatch, V_i is the current value of the i th matrix element and DFTB_i is the value of that element in the initial mio-0-1 parametrization.

Because the results have a strong dependence on λ_{DFTB} , for the remainder of this paper, we will present results in the format of Figure 8. The DFTB penalty is relaxed during

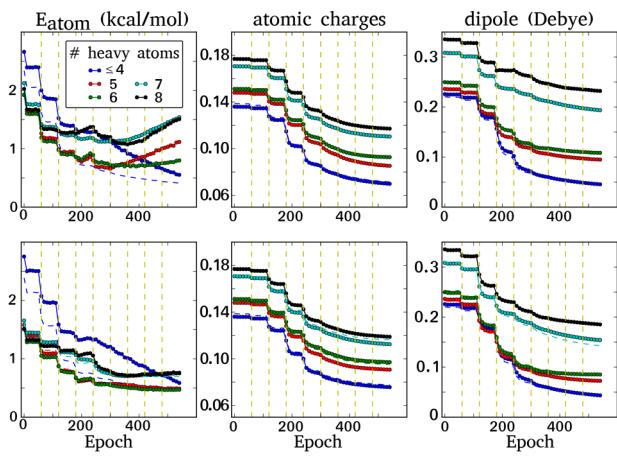


Figure 8. Spline model with regularization used to force monotonic decay, using $\lambda_{\text{mono}} = 10^5$ in eq 29, and to penalize deviation from mio-0-1 values. Dotted lines indicate points at which the penalty for deviations from mio-0-1 starting values are relaxed, with values for λ_{DFTB} following the series: 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, and 10 kcal/mol. Remaining conventions are as in Figure 4.

training by increasing the value of λ_{DFTB} every 60 epochs. The vertical dotted lines in Figure 8 indicate epochs at which the penalty is relaxed. For values of λ_{DFTB} where overfitting is not

prevalent, 60 epochs is sufficient for training to stabilize. This approach therefore provides a concise picture of the dependence of the results on the DFTB penalty, before overfitting sets in, and the manner in which the model degrades, once overfitting sets in. The best performance is at about 360 epochs. Training on up to seven heavy atoms and testing on eight heavy atoms, leads to RMS errors in E_{atom} of 0.72 kcal/mol, q of 0.125, and μ of 0.20 D. Comparison with Table 3 shows that this corresponds to an error reduction of 60% in E_{atom} , 27% in q and 41% in μ . Training on up to four heavy atoms gives similar improvements for E_{atom} (1.08 kcal/mol, 53%) and q (0.12, 27%), with less improvement seen for μ (0.25 D, 28%). The values in parentheses are RMS errors and percent reduction relative to the column in Table 3 for which the reference energy was trained on molecules with up to four heavy atoms.

The Supporting Information shows results from adjusting some classes of matrix elements during the fitting process, freezing the remainder at their mio-1-0 values. The results suggest that most of the improvement results from fitting the off-diagonal matrix elements of \mathbf{H} and \mathbf{G} , and that simultaneous fitting of all matrix elements leads to the best performance.

7. NEURAL NETWORK MODELS

In the spline models of the previous section, the matrix elements depend only on the element type and the distance between atoms. Here, we use neural networks to allow the matrix elements to depend on the molecular environment of the atoms. The features input to the network are those developed for the ANI-1 neural network.²¹ Encoded bonds features¹⁵ were also tried but did not perform as well (see Supporting Information). The form of the network is shown schematically in Figure 9. For each of the three element types

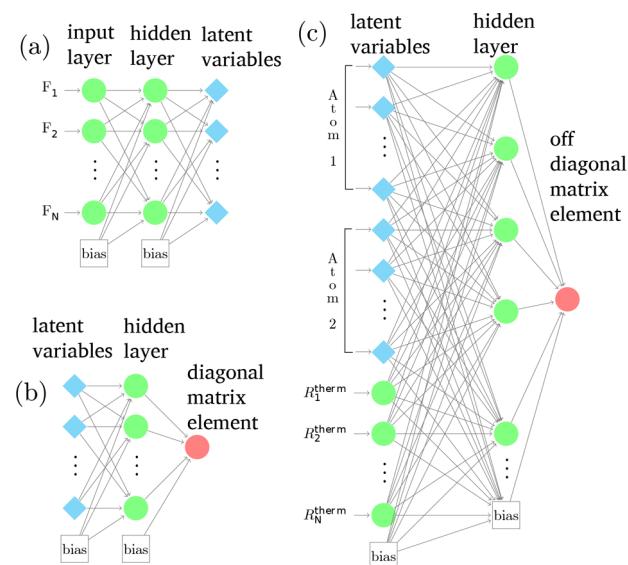


Figure 9. Schematic representation of neural network structure, for the case of one hidden layer. (a) For each element type, C, H, and O, a network takes in features that describe the atomic environment, F_i , and outputs a set of latent variables for that atom. (b) For each of the diagonal matrix element types in Table 1, a neural network predicts the matrix element from the latent variables of the associated atom. (c) Off-diagonal elements are predicted from the latent variables of the two associated atoms and a thermometer-encoded representation of the interatomic distance, eq 31.

(C, H, and O), we create a FFNN that takes the features F_i of a particular atom as input and generates latent variables that summarize the atomic environment of that atom (Figure 9a). This is similar to the use of latent variables in the deep tensor neural network.⁷⁰ Each atom gets a single set of latent variables that are used to predict all matrix elements associated with that atom.

For each of the diagonal matrix element types in Table 1, a FFNN is used to predict the matrix element from the latent variable of the associated atom (Figure 9b). For each of the off-diagonal matrix elements in Table 1, a FFNN is used to predict the matrix element from the latent variables of the two associated atoms, concatenated with features, R_i^{therm} , that thermometer encode the interatomic distance (Figure 9c). The thermometer encoding is done via

$$R_i^{\text{therm}} = \frac{e^{-\sigma(r-r_i^{\text{grid}})}}{1 + e^{-\sigma(r-r_i^{\text{grid}})}} \quad (31)$$

where r is the interatomic distance being encoded, r_i^{grid} is a uniform grid with 10 points, and σ is the inverse of the grid spacing. Similar to the spline model, for interatomic separations beyond the cutoff distances for the Hamiltonian matrix elements in Figure 3, we use the matrix elements from the mio-0-1 parameters. r_i^{grid} is therefore a uniform grid of 10 points spanning the range from the smallest interatomic distances in the data set to the cutoff distance. We emphasize that each atom has a single set of latent variables, that is used in predicting all matrix elements involving that atom.

Although the schematic of Figure 9 shows a single hidden layer for each of the networks, we have tried a number of other network structures (see Supporting Information). Figure 10

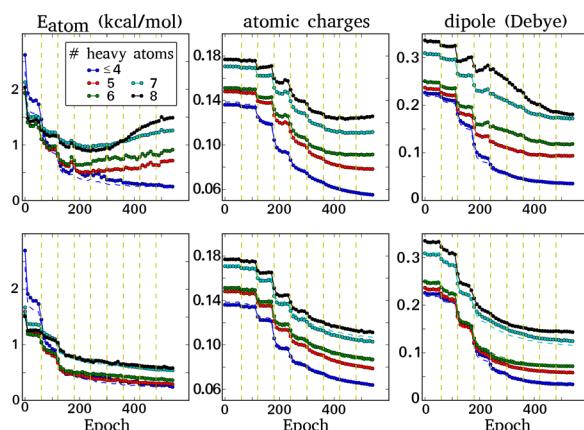


Figure 10. Neural network model with regularization used to penalize deviations from mio-0-1 values. Conventions are as in Figure 8.

shows results from using ANI-1 features,²¹ one hidden layer with 10 nodes for the latent network and the diagonal network, and one hidden layer with 50 nodes for the off-diagonal network. Sigmoid activation was used for all layers, except the final layer which is linear. Initialization occurs in two stages. In the first stage, the weights of each layer are initialized with a random normal distribution, the standard deviation of which is the inverse of the number of inputs to the layer. The distribution is truncated by redrawing any numbers with an absolute value greater than two standard deviations. In the second stage, the full network is then trained to reproduce the mio-0-1 matrix elements for all molecules in the training set for 5000 epochs, leading to an RMS error slightly below 0.1 kcal/mol.

In all cases, the repulsive potential is treated with splines configured as in section 6.

The network is then trained using DFTB regularization in a manner identical to that used while training the spline models. Monotonic regularization is applied only to the repulsive potential. The results of Figure 10 show that the neural network model performs somewhat better than splines. Training on molecules with up to seven heavy atoms and testing on molecules with eight heavy atoms leads to RMS errors, at 500 epochs, in E_{atom} of 0.59 kcal/mol, q of 0.11, and μ of 0.14 D. This corresponds to an error reduction, relative to the spline model in Figure 8, of 18% in E_{atom} , 10% in q , and 27% in μ . When trained on molecules with up to four heavy atoms, the enhancement in performance on molecules with eight heavy atoms, relative to the spline model, is more modest. At 240 epochs, the error in E_{atom} is reduced by 18% relative to the spline, but the errors in q and μ are increased by 21% and 12%, respectively. At 360 epochs, the error in all quantities is reduced relative to the spline model, but by only 1% for E_{atom} , 3% for q , and 9% for μ .

The Supporting Information includes results from a variety of other neural network architectures. These results show that making the model more flexible by adding additional hidden layers or more nodes to the hidden layers, may somewhat improve performance when training on molecules with up to seven heavy atoms, but degrades performance when training on molecules with up to four heavy atoms. Additional approaches to regularization that better restrict the model flexibility may help with transfer from smaller to larger molecules, but developing such regularizations is left to further work.

8. DISCUSSION

The DFTB layer for deep learning introduced here makes it computationally feasible to adjust the electronic portion of a semiempirical QC Hamiltonian to relatively large sets of *ab initio* data. Substantial improvements in predictions of energy, charge distributions, and dipole moments were obtained with both the spline and neural network models explored here. Moreover, because the quantum chemical algorithm is implemented as a layer for deep learning, any model that can be implemented in a deep learning framework may be used to generate the Hamiltonian matrix elements. This opens the possibility of discovering models that further improve performance and expand the range of chemical systems and properties that may be described.

A challenge in development of such models is that the resulting models are highly flexible. This flexibility comes from both the flexibility of models used to generate the Hamiltonian matrix elements and from the large number of such matrix elements present in each molecule. Detrimental effects from this flexibility were apparent in the overfitting that limited transfer of models trained on smaller molecules to larger molecules. This transfer was improved through regularization that penalized deviation of the matrix elements from those of the initial mio-0-1 model parametrization. Finding additional and more effective approaches to regularization is one avenue through which the utility of this approach can potentially be improved.

Despite the highly flexible nature of the model, performance on the training data remained above the 0.5 kcal/mol accuracy for total molecular energy that is the target for chemical accuracy. This indicates that the current model is not able to accurately capture the interactions present in the molecule. The similar performance observed for spline and neural

network models suggests that the limitations are coming from the form of the Hamiltonian itself, as opposed to the models used to generate the matrix elements of this Hamiltonian.

Comparison of DFTB with other semiempirical models suggests one way in which the Hamiltonian can be generalized. In DFTB, the Coulomb interactions between atoms are described by charge–charge interactions. In neglect of diatomic differential overlap (NDDO) models, this is extended to include interactions between atomic dipoles and quadrupoles.²⁷ DFTB has also been extended to include dipole interactions.⁷¹ Extending the current quantum chemical layer to include higher multipole interactions should not lead, as far as we can anticipate, to any fundamental issues.

Another source of information regarding more general forms of the model Hamiltonian comes from comparison with the results from quasiamolecular minimal-basis orbitals (QUAMBO).⁷² QUAMBO can, from the SCF solution of a molecule in a large basis set, generate a minimal basis Hamiltonian that will reproduce that SCF solution. The two-electron integrals of the resulting QUAMBO Hamiltonian are full four-index structures that do not adhere to the highly simplified forms of the DFTB or NDDO Hamiltonians. Inclusion of this aspect into the QC layer seems hopelessly complex. However, examination of QUAMBO Hamiltonians also reveals that the matrix elements between atomic *p* orbitals do not exhibit the cylindrical symmetry implicit in the DFTB model form. This cylindrical symmetry arises from the use of a single pp_{π} model (Table 1), to generate Hamiltonian matrix elements between the *p* orbitals on two atoms. This assumption may be relaxed by allowing models to break the cylindrical symmetry, using features that describe the nonsymmetrical molecular environment about the relevant atoms.

In summary, we hope the DFTB layer developed here will aid development of parametrizations that expand the power and applicability of DFTB. In addition, we hope this will enable more systematic investigations into the benefits and challenges associated with incorporating quantum chemistry directly into deep learning models.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: [10.1021/acs.jctc.8b00873](https://doi.org/10.1021/acs.jctc.8b00873).

Additional details on the chemical data used to fit the models; results from alternative spline and neural net models; and plots of spline models, analogous to those of Figure 5, for all classes of matrix elements ([PDF](#))

AUTHOR INFORMATION

Corresponding Author

*E-mail: yaron@cmu.edu.

ORCID

David J. Yaron: [0000-0001-8485-8685](https://orcid.org/0000-0001-8485-8685)

Funding

This work was supported by the National Science Foundation under Grant CHE-1027985.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors thank Adrian Roitberg for providing early access to the ANI-1 dataset. This work used the Extreme Science and

Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant No. ACI-1548562. Specifically, it used the Bridges system, which is supported by NSF Award No. ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

REFERENCES

- (1) von Lilienfeld, O. A.; Ramakrishnan, R. Machine learning, quantum chemistry, and chemical space. *Rev. Comput. Chem.* **2017**, *30*, 225.
- (2) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **2018**, *361*, 360–365.
- (3) Huan, T. D.; Mannodi-Kanakkithodi, A.; Ramprasad, R. Accelerated Materials Property Predictions and Design Using Motif-based Fingerprints. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2015**, *92*, 014106.
- (4) Meredig, B.; Agrawal, A.; Kirklin, S.; Saal, J. E.; Doak, J.; Thompson, A.; Zhang, K.; Choudhary, A.; Wolverton, C. Combinatorial Screening for New Materials in Unconstrained Composition Space with Machine Learning. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2014**, *89*, 094104.
- (5) OBoyle, N. M.; Campbell, C. M.; Hutchison, G. R. Computational Design and Selection of Optimal Organic Photovoltaic Materials. *J. Phys. Chem. C* **2011**, *115*, 16200–16210.
- (6) Rupp, M.; Tkatchenko, A.; Müller, K.-R.; Von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **2012**, *108*, 058301.
- (7) Ruddigkeit, L.; van Deursen, R.; Blum, L. C.; Reymond, J.-L. Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database GDB-17. *J. Chem. Inf. Model.* **2012**, *52*, 2864–2875.
- (8) Hachmann, J.; Olivares-Amaya, R.; Jinich, A.; Appleton, A. L.; Blood-Forsythe, M. A.; Seress, L. R.; Román-Salgado, C.; Trepte, K.; Atahan-Evrenk, S.; Er, S.; Shrestha, S.; Mondal, R.; Sokolov, A.; Bao, Z.; Aspuru-Guzik, A. Lead Candidates for High-performance Organic Photovoltaics from High-throughput Quantum Chemistry—The Harvard Clean Energy Project. *Energy Environ. Sci.* **2014**, *7*, 698–704.
- (9) Hansen, K.; Montavon, G.; Biegler, F.; Fazli, S.; Rupp, M.; Scheffler, M.; von Lilienfeld, O. A.; Tkatchenko, A.; Müller, K.-R. Assessment and Validation of Machine Learning Methods for Predicting Molecular Atomization Energies. *J. Chem. Theory Comput.* **2013**, *9*, 3404–3419.
- (10) Qu, X.; Latino, D. A.; Aires-de Sousa, J. A Big Data Approach to the Ultra-fast Prediction of DFT-calculated Bond Energies. *J. Cheminf.* **2013**, *5*, 34.
- (11) Sun, J.; Wu, J.; Song, T.; Hu, L.; Shan, K.; Chen, G. Alternative Approach to Chemical Accuracy: A Neural Networks-Based First-Principles Method for Heat of Formation of Molecules Made of H, C, N, O, F, S, and Cl. *J. Phys. Chem. A* **2014**, *118*, 9120–9131.
- (12) Barker, J.; Bulin, J.; Hamaekers, J.; Mathias, S. In *Scientific Computing and Algorithms in Industrial Simulations: Projects and Products of Fraunhofer SCAI*; Griebel, M., Schüller, A., Schweitzer, M. A., Eds.; Springer International Publishing: Cham, 2017; pp 25–42.
- (13) Yao, K.; Herr, J. E.; Brown, S. N.; Parkhill, J. Intrinsic Bond Energies From a Bonds-in-Molecules Neural Network. *J. Phys. Chem. Lett.* **2017**, *8*, 2689–2694.
- (14) Ghiringhelli, L. M.; Vybiral, J.; Levchenko, S. V.; Draxl, C.; Scheffler, M. Big Data of Materials Science: Critical Role of the Descriptor. *Phys. Rev. Lett.* **2015**, *114*, 105503.
- (15) Collins, C. R.; Gordon, G. J.; von Lilienfeld, O. A.; Yaron, D. J. Constant size descriptors for accurate machine learning models of molecular properties. *J. Chem. Phys.* **2018**, *148*, 241718.
- (16) Wang, X.; Wong, L.; Hu, L.; Chan, C.; Su, Z.; Chen, G. Improving the Accuracy of Density-Functional Theory Calculation: The Statistical Correction Approach. *J. Phys. Chem. A* **2004**, *108*, 8514–8525.

(17) Ediz, V.; Monda, A. C.; Brown, R. P.; Yaron, D. J. Using Molecular Similarity to Develop Reliable Models of Chemical Reactions in Complex Environments. *J. Chem. Theory Comput.* **2009**, *5*, 3175–3184.

(18) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Big data meets quantum chemistry approximations: the Δ -machine learning approach. *J. Chem. Theory Comput.* **2015**, *11*, 2087–2096.

(19) Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; von Lilienfeld, O. A. Electronic Spectra from TDDFT and Machine Learning in Chemical Space. *J. Chem. Phys.* **2015**, *143*, 084111.

(20) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1, A data set of 20 million calculated off-equilibrium conformations for organic molecules. *Sci. Data* **2017**, *4*, 170193.

(21) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **2017**, *8*, 3192–3203.

(22) Behler, J.; Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **2007**, *98*, 146401.

(23) Behler, J. Constructing high-dimensional neural network potentials: A tutorial review. *Int. J. Quantum Chem.* **2015**, *115*, 1032–1050.

(24) Boes, J. R.; Groenenboom, M. C.; Keith, J. A.; Kitchin, J. R. Neural network and ReaxFF comparison for Au properties. *Int. J. Quantum Chem.* **2016**, *116*, 979–987.

(25) Shen, L.; Wu, J.; Yang, W. Multiscale quantum mechanics/molecular mechanics simulations with neural networks. *J. Chem. Theory Comput.* **2016**, *12*, 4934–4946.

(26) Artrith, N.; Urban, A. An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for TiO₂. *Comput. Mater. Sci.* **2016**, *114*, 135–150.

(27) Dewar, M. J.; Zoebisch, E. G.; Healy, E. F.; Stewart, J. J. Development and use of quantum mechanical molecular models. 76. AM1: a new general purpose quantum mechanical molecular model. *J. Am. Chem. Soc.* **1985**, *107*, 3902–3909.

(28) Stewart, J. Optimization of parameters for semiempirical methods VI: More modifications to the NDDO approximations and re-optimization of parameters. *J. Mol. Model.* **2013**, *19*, 1–32.

(29) Thiel, W. Semiempirical quantum-chemical methods. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2014**, *4*, 145–157.

(30) Repasky, M. P.; Chandrasekhar, J.; Jorgensen, W. L. PDDG/PM3 and PDDG/MNDO: improved semiempirical methods. *J. Comput. Chem.* **2002**, *23*, 1601–1622.

(31) Cui, Q.; Elstner, M. Density functional tight binding: values of semi-empirical methods in an ab initio era. *Phys. Chem. Chem. Phys.* **2014**, *16*, 14368–14377.

(32) Welborn, M.; Cheng, L.; Miller, T. F. III Transferability in Machine Learning for Electronic Structure via the Molecular Orbital Basis. *J. Chem. Theory Comput.* **2018**, *14*, 4772–4779.

(33) Girshick, R. Fast r-cnn. *Proc. IEEE Intl. Conf. Comput. Vision* **2015**, 1440–1448.

(34) Ren, S.; He, K.; Girshick, R.; Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems* **2015**, *39*, 91–99.

(35) Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. *Proc. IEEE Conf. Comput. Vision Pattern Recognit.* **2016**, 779–788.

(36) Collobert, R.; Weston, J. A unified architecture for natural language processing: Deep neural networks with multitask learning. *Proc. 25th Int'l Conf. Mach. Learnning* **2008**, 160–167.

(37) Collobert, R.; Weston, J.; Bottou, L.; Karlen, M.; Kavukcuoglu, K.; Kuksa, P. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **2011**, *12*, 2493–2537.

(38) Elstner, M.; Porezag, D.; Jungnickel, G.; Elsner, J.; Haugk, M.; Frauenheim, T.; Suhai, S.; Seifert, G. Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1998**, *58*, 7260.

(39) Koskinen, P.; Mäkinen, V. Density-functional tight-binding for beginners. *Comput. Mater. Sci.* **2009**, *47*, 237–253.

(40) Nishizawa, H.; Nishimura, Y.; Kobayashi, M.; Irle, S.; Nakai, H. Three pillars for achieving quantum mechanical molecular dynamics simulations of huge systems: Divide-and-conquer, density-functional tight-binding, and massively parallel computation. *J. Comput. Chem.* **2016**, *37*, 1983–1992.

(41) Sattelmeyer, K. W.; Tirado-Rives, J.; Jorgensen, W. L. Comparison of SCC-DFTB and NDDO-based semiempirical molecular orbital methods for organic molecules. *J. Phys. Chem. A* **2006**, *110*, 13551–13559.

(42) Yang, Y.; Yu, H.; York, D.; Elstner, M.; Cui, Q. Description of phosphate hydrolysis reactions with the self-consistent-charge density-functional-tight-binding (SCC-DFTB) theory. 1. Parameterization. *J. Chem. Theory Comput.* **2008**, *4*, 2067–2084.

(43) Dolgonos, G.; Aradi, B.; Moreira, N. H.; Frauenheim, T. An improved self-consistent-charge density-functional tight-binding (SCC-DFTB) set of parameters for simulation of bulk and molecular systems involving titanium. *J. Chem. Theory Comput.* **2009**, *6*, 266–278.

(44) Wahiduzzaman, M.; Oliveira, A. F.; Philipsen, P.; Zhechkov, L.; van Lenthe, E.; Witek, H. A.; Heine, T. DFTB parameters for the periodic table: Part 1, Electronic structure. *J. Chem. Theory Comput.* **2013**, *9*, 4006–4017.

(45) Oliveira, A. F.; Philipsen, P.; Heine, T. DFTB parameters for the periodic table, part 2: energies and energy gradients from hydrogen to calcium. *J. Chem. Theory Comput.* **2015**, *11*, 5209–5218.

(46) Zheng, G.; Witek, H. A.; Bobadova-Parvanova, P.; Irle, S.; Musaev, D. G.; Prabhakar, R.; Morokuma, K.; Lundberg, M.; Elstner, M.; Köhler, C.; Frauenheim, T. Parameter calibration of transition-metal elements for the spin-polarized self-consistent-charge density-functional tight-binding (DFTB) method: Sc, Ti, Fe, Co, and Ni. *J. Chem. Theory Comput.* **2007**, *3*, 1349–1367.

(47) Knaup, J. M.; Hourahine, B.; Frauenheim, T. Initial Steps toward Automating the Fitting of DFTB Erep(r). *J. Phys. Chem. A* **2007**, *111*, S637–S641.

(48) Gaus, M.; Chou, C.-P.; Witek, H.; Elstner, M. Automatized parametrization of SCC-DFTB repulsive potentials: Application to hydrocarbons. *J. Phys. Chem. A* **2009**, *113*, 11866–11881.

(49) Bodrog, Z.; Aradi, B.; Frauenheim, T. Automated repulsive parametrization for the DFTB method. *J. Chem. Theory Comput.* **2011**, *7*, 2654–2664.

(50) Lourenço, M. P.; da Silva, M. C.; Oliveira, A. F.; Quintão, M. C.; Duarte, H. A. FASP: a framework for automation of Slater–Koster file parameterization. *Theor. Chem. Acc.* **2016**, *135*, 250.

(51) Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; Kudlur, M.; Levenberg, J.; Monga, R.; Moore, S.; Murray, D. G.; Steiner, B.; Tucker, P.; Vasudevan, V.; Warden, P.; Wicke, M.; Yu, Y.; Zheng, X. Tensorflow: a system for large-scale machine learning. *OSDI 2016*, 265–283.

(52) Nystrom, N. A.; Levine, M. J.; Roskies, R. Z.; Scott, J. R. Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics. *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*. New York, NY, USA, 2015; pp 30:1–30:8.

(53) Kranz, J. J.; Kubillus, M.; Ramakrishnan, R.; von Lilienfeld, O. A.; Elstner, M. Generalized Density-Functional Tight-Binding Repulsive Potentials from Unsupervised Machine Learning. *J. Chem. Theory Comput.* **2018**, *14*, 2341–2352.

(54) Chou, C.-P.; Nishimura, Y.; Fan, C.-C.; Mazur, G.; Irle, S.; Witek, H. A. Automatized Parameterization of DFTB Using Particle Swarm Optimization. *J. Chem. Theory Comput.* **2016**, *12*, 53–64.

(55) Krishnapriyan, A.; Yang, P.; Niklasson, A. M. N.; Cawkwell, M. J. Numerical Optimization of Density Functional Tight Binding Models: Application to Molecules Containing Carbon, Hydrogen, Nitrogen, and Oxygen. *J. Chem. Theory Comput.* **2017**, *13*, 6191–6200.

(56) Kullgren, J.; Wolf, M. J.; Hermansson, K.; Khler, C.; Aradi, B.; Frauenheim, T.; Broqvist, P. Self-Consistent-Charge Density-Func-

tional Tight-Binding (SCC-DFTB) Parameters for Ceria in 0D to 3D. *J. Phys. Chem. C* **2017**, *121*, 4593–4607.

(57) Li, Y.; Qi, Y. Transferable Self-Consistent Charge Density Functional Tight-Binding Parameters for LiMetal and Li-Ions in Inorganic Compounds and Organic Solvents. *J. Phys. Chem. C* **2018**, *122*, 10755–10764.

(58) Huran, A. W.; Steigemann, C.; Frauenheim, T.; Aradi, B.; Marques, M. A. L. Efficient Automatized Density-Functional Tight-Binding Parametrizations: Application to Group IV Elements. *J. Chem. Theory Comput.* **2018**, *14*, 2947–2954.

(59) Weber, W.; Thiel, W. Orthogonalization corrections for semiempirical methods. *Theor. Chem. Acc.* **2000**, *103*, 495–506.

(60) Dral, P. O.; von Lilienfeld, O. A.; Thiel, W. Machine learning of parameters for accurate semiempirical quantum chemical calculations. *J. Chem. Theory Comput.* **2015**, *11*, 2120–2125.

(61) Schütt, O.; VandeVondele, J. Machine learning adaptive basis sets for efficient large scale DFT simulation. *J. Chem. Theory Comput.* **2018**, *14*, 4168.

(62) Kinga, D.; Adam, J. B. A method for stochastic optimization. *Int. Conf. Learn. Representations*. **2015**.

(63) Pulay, P. *J. Comput. Chem.* **1982**, *3*, 556–560.

(64) Frisch, M. J.; Trucks, G. W.; Schlegel, H. B.; Scuseria, G. E.; Robb, M. A.; Cheeseman, J. R.; Scalmani, G.; Barone, V.; Mennucci, B.; Petersson, G. A.; Nakatsuji, H.; Caricato, M.; Li, X.; Hratchian, H. P.; Izmaylov, A. F.; Bloino, J.; Zheng, G.; Sonnenberg, J. L.; Hada, M.; Ehara, M.; Toyota, K.; Fukuda, R.; Hasegawa, J.; Ishida, M.; Nakajima, T.; Honda, Y.; Kitao, O.; Nakai, H.; Vreven, T.; Montgomery, J. A., Jr.; Peralta, J. E.; Ogliaro, F.; Bearpark, M.; Heyd, J. J.; Brothers, E.; Kudin, K. N.; Staroverov, V. N.; Kobayashi, R.; Normand, J.; Raghavachari, K.; Rendell, A.; Burant, J. C.; Iyengar, S. S.; Tomasi, J.; Cossi, M.; Rega, N.; Millam, J. M.; Klene, M.; Knox, J. E.; Cross, J. B.; Bakken, V.; Adamo, C.; Jaramillo, J.; Gomperts, R.; Stratmann, R. E.; Yazyev, O.; Austin, A. J.; Cammi, R.; Pomelli, C.; Ochterski, J. W.; Martin, R. L.; Morokuma, K.; Zakrzewski, V. G.; Voth, G. A.; Salvador, P.; Dannenberg, J. J.; Dapprich, S.; Daniels, A. D.; Farkas, O.; Foresman, J. B.; Ortiz, J. V.; Cioslowski, J.; Fox, D. J. *Gaussian 09*; revision D.01; Gaussian Inc.: Wallingford CT, 2009.

(65) Hu, H.; Lu, Z.; Yang, W. Fitting molecular electrostatic potentials from quantum mechanical calculations. *J. Chem. Theory Comput.* **2007**, *3*, 1004–1013.

(66) Krüger, T.; Elstner, M.; Schiffels, P.; Frauenheim, T. Validation of the density-functional based tight-binding approximation method for the calculation of reaction energies and other data. *J. Chem. Phys.* **2005**, *122*, 114110.

(67) Aradi, B.; Hourahine, B.; Frauenheim, T. DFTB+, a sparse matrix-based implementation of the DFTB method. *J. Phys. Chem. A* **2007**, *111*, 5678–5684.

(68) Niehaus, T. A.; Elstner, M.; Frauenheim, T.; Suhai, S. Application of an approximate density-functional method to sulfur containing compounds. *J. Mol. Struct.: THEOCHEM* **2001**, *541*, 185–194.

(69) De Boor, C.; De Boor, C.; Mathématicien, E.-U.; De Boor, C.; De Boor, C. *A Practical Guide to Splines*; Springer-Verlag New York, 1978; Vol. 27.

(70) Schütt, K. T.; Arbabzadah, F.; Chmiela, S.; Müller, K. R.; Tkatchenko, A. Quantum-chemical insights from deep tensor neural networks. *Nat. Commun.* **2017**, *8*, 13890.

(71) Wu, Y.; Ilie, A.; Crampin, S. Self-consistent charge and dipole density functional tight binding method and application to carbon-based systems. *Comput. Mater. Sci.* **2017**, *134*, 206–213.

(72) Lu, W. C.; Wang, C. Z.; Schmidt, M. W.; Bytautas, L.; Ho, K. M.; Ruedenberg, K. Molecule intrinsic minimal basis sets. I. Exact resolution of ab initio optimized molecular orbitals in terms of deformed atomic minimal-basis orbitals. *J. Chem. Phys.* **2004**, *120*, 2629–2637.