

IMPROVING NEURAL NETWORK PREDICTIONS OF MATERIAL PROPERTIES WITH LIMITED DATA USING TRANSFER LEARNING

*Schuyler Krawczuk & Daniele Venturi**

Department of Applied Mathematics, University of California Santa Cruz, Santa Cruz, CA 95064, USA

*Address all correspondence to: Daniele Venturi, Department of Applied Mathematics, University of California Santa Cruz, Santa Cruz, CA 95064, USA,
E-mail: venturi@ucsc.edu

Original Manuscript Submitted: 6/29/2020; Final Draft Received: 10/22/2020

We develop new transfer learning algorithms to accelerate prediction of material properties from ab initio simulations based on density functional theory. Transfer learning has been successfully utilized for data-efficient modeling in applications other than materials science, and it allows transferable representations learned from large datasets to be repurposed for learning new tasks even with small datasets. In the context of materials science, this opens the possibility to develop generalizable neural network models that can be repurposed on other materials, without the need of generating a large (computationally expensive) training set of materials properties. The proposed transfer learning algorithms are demonstrated on predicting the Gibbs free energy of light transition metal oxides.

KEY WORDS: *density functional theory, transfer learning, materials science*

1. INTRODUCTION

Determining physical properties of materials using computer simulations is commonplace in many areas of science, including chemistry, physics, and engineering. This usually relies on *ab initio* calculations, which are based around solving the Schrödinger equation (Zinn-Justin, 2002) for these materials. The Schrödinger equation governs the wave function of a many-body quantum system, from which information such as its electronic structure can be found, and in turn used to determine useful properties. Directly solving the Schrödinger equation for non-trivial systems quickly becomes computationally infeasible as the number of atoms increases. This led to the rise of methods such as density functional theory (DFT) (Parr and Weitao, 1994), which bypasses the need for explicitly solving for the wave function to find the electronic structure, thus opening the way to perform practical computations.

The state space of a quantum chemical system for even a narrow application of materials can be very large, as both content of elements and atomic position may be taken into account. To thoroughly explore this space can take significant computational time using DFT, so there is a demand to further expedite property prediction beyond current state-of-the-art capabilities. A leading method for this is machine learning. Given sufficient examples of materials, a neural

network can be trained to represent the mapping between atomic structure and material properties such that it can generalize well to similar but unseen materials. The benefit of this approach is that once trained, a neural network can make predictions in a matter of milliseconds rather than the minutes or hours DFT can take for a single material. Several neural network architectures have been developed specifically for modeling atomistic systems (Behler and Parrinello, 2007; Schütt et al., 2017a,b), leading to increasingly accurate representations. A caveat of current neural network approaches is that it can take a large number of examples to learn a sufficiently accurate mapping. Although there are large public databases available with material structures and some of their properties (Curtarolo et al., 2012; Jain et al., 2013), they may not sufficiently cover the target material space, leading to the need of generating more training data. In these cases, it is desirable to minimize this extra cost associated with generating additional data, and training accurate neural networks with smaller datasets.

This can be achieved by transfer learning, which is a popular method used for training neural networks with small datasets (Pan and Yang, 2010). Transfer learning reuses representations learned by neural networks trained on large datasets. Instead of randomly initializing weights for training on a certain task, the weights of a neural network trained on a similar task are used as a starting point to provide a more optimal initialization. Transfer learning has been very successful in the field of computer vision, where large image datasets are publicly available, for which pre-trained models are commonly made available (Razavian et al., 2014). These large datasets led to robust machine learning and generalizable features that are also useful for other tasks (Pan and Yang, 2010). Similarly, in material science, large databases of material properties have been made publicly available (Razavian et al., 2014). These databases can be leveraged for transfer learning, allowing pre-trained neural networks to be used for initialization without the creation of new data. For instance, Yamada et al. (2019) show transfer learning has been successfully adopted for increasing correlation in material property predictions given a very small number of samples. Transfer learning also has been used for correcting DFT predictions on benchmarks for reaction thermochemistry, isomerization, and drug-like molecular torsions (Smith et al., 2019). Along the same line of research as Smith et al. (2019), in this paper we develop transfer learning algorithms to accelerate prediction of material properties from *ab initio* simulations based on DFT.

This paper is organized as follows. In Section 2 we provide a brief review of *ab initio* calculation of material properties using DFT. In Section 3 we discuss continuous-filter convolutional neural network representations of atomistic systems (SchNet) (Schütt et al., 2017b) and their training using the ADAM algorithm (Kingma and Ba, 2017). In Section 4 we develop transfer learning schemes that employ SchNet to improve training of neural network models with small datasets. The accuracy of these methods is investigated in Section 5 in application problems involving transition metal oxides. The main findings are summarized in Section 6.

2. BRIEF REVIEW OF DENSITY FUNCTIONAL THEORY

Ab initio calculations for electronic structure rely on solving the time-independent Schrödinger equation. The full state of the system is computed by solving an eigenvalue problem of the form

$$H\Psi = E\Psi, \quad (1)$$

where H is the Hamiltonian operator of the system, consisting of the sum of kinetic (T) and potential energies (V), while E represents the energy of a specific state (eigenvalue) described by

the wave function Ψ . In atomic structure calculations, nuclei can be treated as a static external potential V_{ext} , and only the electrons are considered in the wave equation (Born and Oppenheimer, 1927). Hence, the wave function Ψ in Eq. (1) for N electrons in a three-dimensional Euclidean space has $3N$ degrees of freedom

$$\Psi = \Psi(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N). \quad (2)$$

Additionally, each electron has an interaction with each other. As the number of electrons increases, not only does the number of dimensions of the wave function increase, but the number of electron-electron interaction terms in the Hamiltonian increases exponentially (Martin, 2004). In general, the many-body Hamiltonian can be written as[†]

$$H = T + V_{\text{ext}} + V_{\text{e-e}} \quad (3)$$

$$= \sum_{i=1}^N \left[\frac{-\hbar^2}{2m_e} \nabla_{\mathbf{r}_i}^2 + V_{\text{ext}}(\mathbf{r}_i) + \sum_{j>i}^N \frac{e^2}{|\mathbf{r}_i - \mathbf{r}_j|} \right]. \quad (4)$$

Despite recent advances in high-dimensional approximation theory (Cho et al., 2017; Dektor and Venturi, 2020; Han et al., 2020; Khoromskaia and Khoromskij, 2018), solving the eigenvalue problem [Eq. (1)] for a non-trivial system involving many atoms quickly becomes computationally infeasible.

DFT (Parr and Weitao, 1994) was originally proposed to mitigate such dimensionality problems. The key idea relies on expressing material properties (ground state) as a functional of the charge density $n(\mathbf{r})$ rather than relying on the wave function Ψ . The charge density is only three-dimensional regardless of the number of electrons, leading to calculations that scale with N more efficiently. The fact that ground state material properties can be expressed as a functional of the charge density was proved in two celebrated theorems by Hohenberg and Kohn (1964). For instance, total ground state energy of an atomistic system can be written as

$$E[n] = F[n] + \int_{\mathbb{R}^3} V_{\text{ext}}(\mathbf{r})n(\mathbf{r})d\mathbf{r}, \quad (5)$$

where $F[n]$ is some functional of n (Venturi, 2018; Venturi and Dektor, 2020), and the second term on the right-hand side represents the interaction of electrons with the external potential created from the nuclei. The Hohenberg–Kohn theorems showed that the functional $F[n]$ exists. However they give no guidance on how to find it.

The Kohn–Sham equations were later introduced, which put the Hohenberg–Kohn findings to use and gave a practical approach to finding this functional. The Kohn–Sham theory relies on the assumption that a non-interacting system of electrons will have the same electron density as an interacting system of the same structure. Based on this, a Schrödinger-like equation can be solved for each individual electron, and all resulting wave functions, called the Kohn–Sham orbitals, can be used to calculate the electron density. In this setting, the eigenvalue problem [Eq. (1)] is replaced by a simpler eigenvalue problem of the form

$$H_{\text{eff}}\psi_i = \epsilon_i\psi_i, \quad (6)$$

where H_{eff} is an *effective* Hamiltonian for this fictitious system. The effective Hamiltonian can be written as

$$H_{\text{eff}} = T[n] + V_{\text{Hartree}}[n] + V_{\text{XC}}[n] + \int_{\mathbb{R}^3} V_{\text{ext}}(\mathbf{r})n(\mathbf{r})d\mathbf{r}. \quad (7)$$

[†]In Eq. (3) $V_{\text{e-e}}$ denotes the electron-electron interaction potential.

This first term at the right-hand side represents the kinetic energy of this system. The second term, $V_{\text{Hartree}}[n]$, is the Hartree potential, which accounts for the repulsion between the electrons:

$$V_{\text{Hartree}}[n] = \frac{1}{2} \int \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'. \quad (8)$$

The third term in Eq. (7), $V_{\text{XC}}[n]$, is known as exchange-correlation potential, and it approximates more complicated interactions between the electrons. The exchange-correlation functional's exact form is unknown, but it is also the smallest contribution to the total energy. Because of this, it can be approximated and can still lead to an accurate solution. The simplest approximation used for exchange-correlation is the local-density approximation (LDA) (Parr and Weitao, 1994). As stated by Kohn and Sham, solids can be considered close to the limit of the uniform electron gas (Hohenberg and Kohn, 1964). The local exchange-correlation energy for the uniform electron gas is known, written as ϵ_{xc} . The LDA exchange-correlation function can be written as

$$E_{\text{XC}}[n] = \int n(\mathbf{r}) \epsilon_{\text{xc}}(n(\mathbf{r})) d\mathbf{r}. \quad (9)$$

A more accurate class of functionals that build off of LDA is the generalized-gradient approximation (GGA). These are exchange-correlation functionals that include a term F_{XC} that is in terms of the gradient of the density; that is,

$$E_{\text{XC}}[n] = \int n(\mathbf{r}) \epsilon_{\text{xc}}(n(\mathbf{r})) F_{\text{XC}}(\nabla n(\mathbf{r})) d\mathbf{r}. \quad (10)$$

The significant improvement in accuracy given by GGA functionals led to the wider adoption of DFT across chemistry and material science (Martin, 2004). The exchange-correlation potential V_{XC} in Eq. (7) is the first-order functional derivative (Venturi and Dektor, 2020) of $E_{\text{XC}}[n]$:

$$V_{\text{XC}}[n] = \frac{\delta E_{\text{XC}}[n]}{\delta n}. \quad (11)$$

To solve the eigenvalue problem [Eq. (6)] the Kohn–Sham orbitals ψ_i are usually represented relative to a finite-dimensional basis $\{\phi_1, \dots, \phi_q\}$ as

$$\psi_i = \sum_{j=1}^q c_{ij} \phi_j. \quad (12)$$

A substitution of Eq. (12) into Eq. (6) and subsequent projection onto the basis $\{\phi_k\}$ yields a generalized eigenvalue problem for c_{ij} , which is typically solved with an iterative method (Blaha et al., 2010).

Once the Kohn–Sham orbital ψ_i is computed for each electron in the system, the electron density can be obtained as

$$n(\mathbf{r}) = \sum_{i=1}^N |\psi_i(\mathbf{r})|^2. \quad (13)$$

This density can then be used to obtain ground state properties; for instance, the total energy function

$$E[n] = T[n] + E_{\text{Hartree}}[n] + E_{\text{XC}}[n] + \int_{\mathbb{R}^3} V_{\text{ext}}(\mathbf{r}) n(\mathbf{r}) d\mathbf{r}. \quad (14)$$

The effective Hamiltonian [Eq. (7)] requires a charge density to construct it initially, meaning an iterative approach needs to be taken to finding $n(\mathbf{r})$. If the n used to construct the effective Hamiltonian is consistent with the n that results from solving the Kohn–Sham equations, that charge density is consistent for the system. A self-consistent loop, seen in Fig. 1, is followed until consistency is reached. In Fig. 2 we plot the results of a three-dimensional DFT calculation for the electron density of a silicon compound.

Computational cost can be a limiting factor in high-throughput calculations with DFT. To give a sense of this, 1300 calculations using the algorithm described above were done for transition metal oxides in this work. These materials had an average of 30 atoms each. Calculations were done on a compute node with two Intel 2.1 GHz Xeon E5-2620v4 processors (64 GB RAM) using 16 cores. As shown in Fig. 3, the mean calculation time for one run was around 27 minutes, with the quickest taking 5 seconds and the longest taking 20 hours.

3. NEURAL NETWORK REPRESENTATION OF ATOMISTIC SYSTEMS

Neural networks are a natural choice for surrogate modeling of DFT. Their ability to learn the nonlinear mapping from an atomic structure to a property has the benefit of not relying on the bias hand-picked features for input, instead solely deriving relationships from the unprocessed data, as visualized in Fig. 4. However, there are some constraints that must be imposed upon the neural network architecture to help enforce the underlying physics. For instance, different

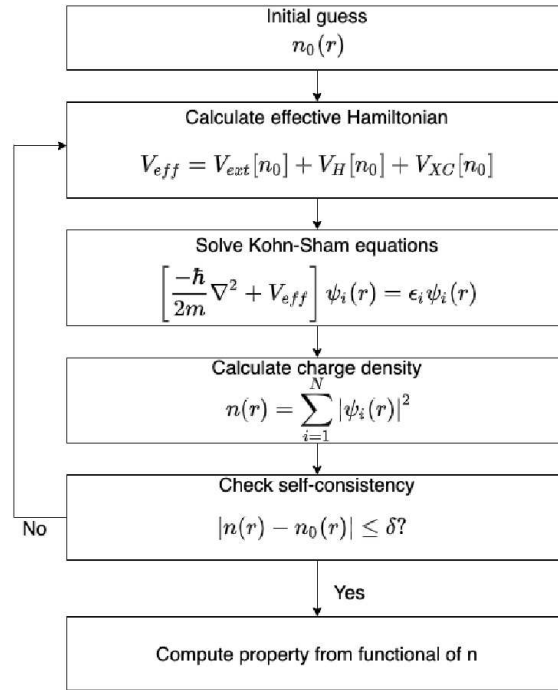


FIG. 1: Self-consistent loop for solving the Kohn–Sham equations. First, a guess for the electron density $n(\mathbf{r})$ is created based on atomic structure. Using this $n(\mathbf{r})$, the Kohn–Sham equations are solved and the resulting $n(\mathbf{r})$ is compared to the initial guess. If these are not within a tolerance (δ) of each other, n is updated as a linear combination of the initial and final and another iteration is performed.




FIG. 2: Top: From top left moving clockwise: plot of V_{eff} , V_{ext} , V_{XC} , and V_{Hartree} in a two-dimensional slice of a Si_2 crystal, along the same plane as the two atoms in the unit cell. Spatial units are in Bohrs (5.29×10^{-11} m) and potential is in Rydberg units (13.6 eV). Bottom: The resulting electron density (in electrons per cubic Bohr) for the above potential.




FIG. 3: Histogram of DFT calculation times for 1300 transition metal oxides. DFT simulations were performed on a compute node with two Intel 2.1 GHz Xeon E5-2620v4 processors (64 GB RAM) using 16 cores.

atomistic systems can vary in the number of atoms, while a neural network typically works with a fixed input size. Also, the result should be invariant to the order atoms are included in the input, as well as the orientation of the entire molecule. These characteristics are not naturally

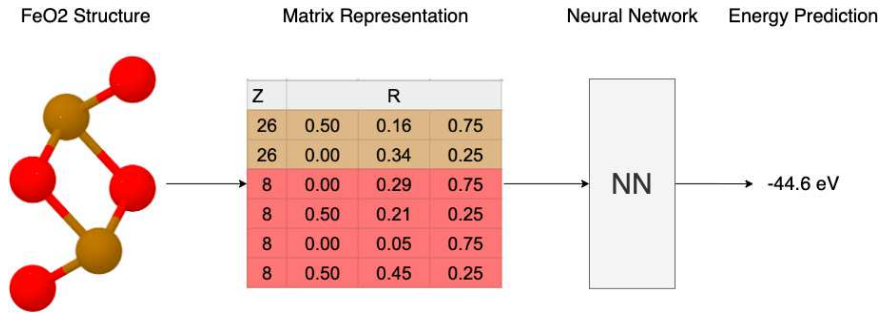


FIG. 4: Atomistic neural network approach to predicting material properties. The structure of a material can be represented by a matrix consisting of each atom's atomic number (Z) and its Cartesian coordinates (R) with respect to a reference point. This representation will be used as the input to the neural network, which maps it to a target property, such as the exchange-correlation potential of the total energy of the compound.

accounted for in typical feed-forward neural networks. A significant research effort indeed has been recently focused in addressing these issues, with specific reference to atomistic systems.

In particular, Schütt et al. (2017b) proposed a continuous-filter convolutional neural network called SchNet for modeling quantum interactions. In our paper we utilize such framework for both its nearly state-of-the-art performance on benchmark datasets and the SchNetPack toolbox released with it, which provides a simple framework for working with atomistic systems. SchNet scales well with variably-sized inputs by using the same weights for each atom in an atomic system, resulting in per-atom contributions. Additionally, SchNet employs continuous-filter convolutional layers. Convolutional layers are the state-of-the-art tools for machine learning with spatial data, but typically these are discretized, such as pixels of an image. Molecular structure does not lie on a grid such as these signals. Although it can be discretized, it requires choosing a proper interpolation scheme and typically a large number of grid points for proper representation that can capture subtle positional changes of atoms. Continuous-filter convolutional layers are implemented in SchNet, getting around this problem by applying a convolution element-wise. Given feature representations of n objects $\mathbf{X}^l = (\mathbf{x}_1^l, \dots, \mathbf{x}_n^l)$ at locations $\mathbf{R}^l = (\mathbf{r}_1^l, \dots, \mathbf{r}_n^l)$, the output of the continuous convolutional layer l at position \mathbf{r}_i is

$$\mathbf{x}_i^{l+1} = (\mathbf{X}^l * \mathbf{W}^l)_i = \sum_{j=1}^n \mathbf{x}_j^l \odot \mathbf{W}^l(\mathbf{r}_i - \mathbf{r}_j), \quad (15)$$

where $*$ denotes the convolution operator, and \odot is the element-wise (Hadamard) product. The filter \mathbf{W}^l weights the distance between the atoms in the system. In the continuous-filter convolutional layer, the distances $d_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ are expanded relative to radial basis functions as

$$e_k(d_{ij}) = \exp(-\gamma|d_{ij} - \mu_k|^2), \quad (16)$$

located at centers $0 \text{ \AA} \leq \mu_k \leq 30 \text{ \AA}$ with $\gamma = 10 \text{ \AA}$. Introducing this additional nonlinearity causes the filter to be less correlated, since the network after initialization is close to linear. This speeds up the beginning of the training process, which may plateau otherwise (Schütt et al., 2017b). In addition to the radial basis functions, each atom is represented by an embedding unique to its atomic number. This embedding is a vector of a predefined length F that is refined through each layer l of the network. The feature representations for an n -atom system for a layer

are $X^l = (x_1^l, \dots, x_n^l)$, with $x_i^l \in \mathbb{R}^F$. Each feature vector is initialized randomly for each Z_i such that

$$x_i^0 = a_{Z_i}, \quad (17)$$

and it is refined during training.

The SchNet architecture consists of the previously described features in blocks called the interaction blocks. Each interaction block refines the feature representations, which are then passed to a final set of atom-wise layers and are pooled to reach the output value. Figure 5 shows the full form of this output. Rather than each interaction block being a composition of the previous, as typically done with neural network layers, each uses a residual connection. The features are updated in each layer as

$$x_i^{l+1} = x_i^l + v_i^l. \quad (18)$$

This connection helps to prevent overfitting, as it is easier for v_i^l to become zero in the training process if the next layer is unnecessary by minimizing the residual between x_i^{l+1} and x_i^l . Without a residual connection, the next layer would be updated as

$$x_i^{l+1} = f^l(x_i^l), \quad (19)$$

requiring f^l to learn the identity function, which is a non-trivial task. The activation function used is the shifted softplus, which is defined as

$$\text{ssp}(x) = \ln(0.5e^x + 0.5). \quad (20)$$

This function is a smooth approximation of the ReLu. The atom-wise layer in the SchNet architecture shown in Fig. 5 applies an affine transformation to the features of each atom separately. This layer shares the same weights throughout every atom, giving the output

$$x_n^{l+1} = W^l x_i^l + b^l, \quad (21)$$

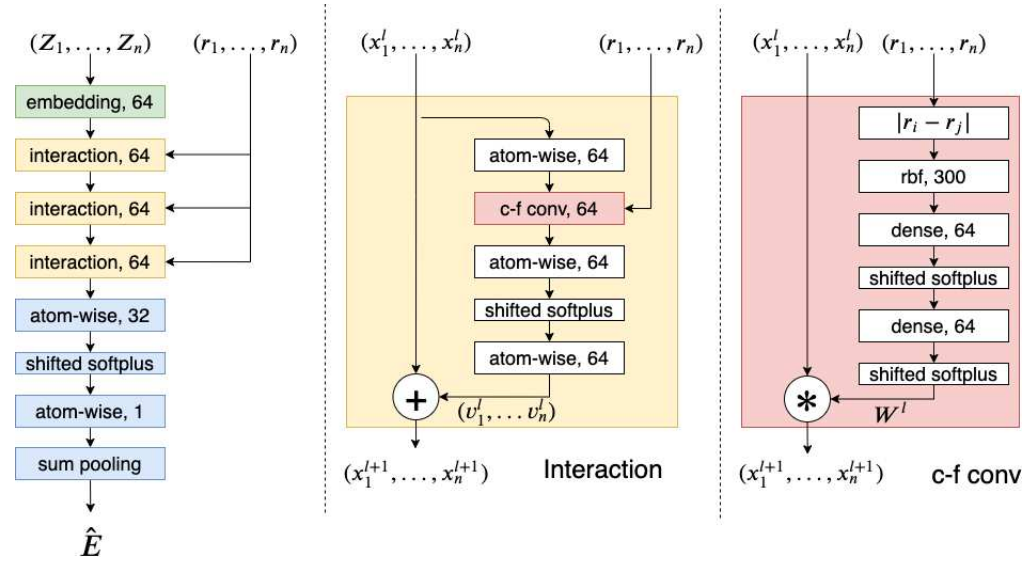


FIG. 5: SchNet architecture using a feature size of 64 and three interaction blocks. The interaction block is shown in the middle, and the continuous-filter convolutional layer on the right.

for the atom i . The sharing of weights across all atoms allows for the network to scale with the size of the system properly.

3.1 Training the SchNet Architecture

To train SchNet in a supervised learning setting we minimize the mean square error between the predicted property and its observed value using the adaptive moment (ADAM) algorithm. ADAM is a variation of the stochastic gradient descent. In classical stochastic gradient descent the weights of the neural network β are updated based on the gradient of the cost function E as

$$\beta_{k+1} = \beta_k - \alpha \nabla E, \quad (22)$$

for step k and learning rate α . It is possible to improve the convergence rate of stochastic gradient descent by multiplying the learning rate by a factor of the previous iteration's step. In this setting, for each iteration, β is updated as

$$\beta_{k+1} = \beta_k + v, \quad (23)$$

where

$$v = \eta v - \alpha g, \quad g = \nabla E. \quad (24)$$

In Eq. (24) $\eta \in [0, 1)$ is a predetermined parameter (Goodfellow et al., 2016), and v must be initialized. The algorithm of Eqs. (23) and (24) is also known as gradient descent with momentum, in a physical analogy for the velocity of a ball rolling down a hill. In fact, as the ball rolls down the hill, much like the optimization descends toward a minimum, the ball will gain speed. The larger η is, the more the previous iteration will affect the next one. This momentum also helps to escape local minima, as seen in Fig. 6.

The ADAM algorithm updates individual learning rates based on previous learning rates in an exponential moving average (Kingma and Ba, 2017). Each parameter is updated similarly to Eqs. (23) and (24). Specifically, in ADAM we have that Eq. (24) is replaced by

$$v = \eta v - \frac{\alpha}{\sqrt{r}} \odot g, \quad (25)$$

where r (accumulation of the gradient), is updated as

$$r = \rho r + (1 - \rho) g \odot g. \quad (26)$$




FIG. 6: One-dimensional function minimization using gradient descent with and without momentum. Gradient descent with momentum gradually increases the step size during the descent, allowing it to escape the local minimum at $x = 5$, to which gradient descent converges.

Although ADAM exhibits faster convergence to a minimum, it has been shown to not reach an optimal solution as well as regular stochastic gradient descent, which is more likely to reach a value close to the global minimum (Reddi et al., 2019; Wilson et al., 2017). Similarly to the training scheme used with SchNetPack’s released models, an initial learning rate α of 10^{-4} is used in this paper. For each training plateau, where the training loss does not decrease for 25 iterations, the learning rate is reduced by a factor of 0.8 until a minimum learning rate of 10^{-6} is reached.

4. TRANSFER LEARNING WITH SCHNET

In this section we develop a transfer learning scheme in conjunction with SchNet to improve the ability of models to generalize when trained with small training sets. Transfer learning reuses representations learned in training with a source task as a starting point for training on a different but similar target task. In this case, a SchNet model is trained with a large source dataset, and the optimized weights are reused for initializing training with the target dataset.

4.1 Transition Metal Oxide Database

In this paper, transition metal oxides are considered for property prediction. Transition metal oxides are used in practical applications for solar energy conversion. However, poor conductivity and electron-hole separation limits their carrier conductivity. It has been shown that appropriate doping (adding of impurity) of these materials may improve their utility. An important property to be found in these doped transition metal oxides is a low defect formation energy. The defect formation energy is the difference between the total free energy of the pure transition metal oxide and that of the impure, doped transition metal oxide. Although there are other properties of importance, this work focuses on the learning of the mapping between transition metal oxides and their free energy.

The target dataset we first considered is composed of 517 transition metal oxides of iron, titanium, and vanadium, along with the corresponding free energies of the compounds. These materials each have between 2 and 110 atoms, with a mean of 12. An extension of this dataset was then created for further testing of the transfer learning methods, introducing an additional 146 transition metal oxides consisting of chromium and manganese. The source dataset is a subset of the Materials Project database (Jain et al., 2013). This subset consisted of materials made up of 87 elements, including those of the target dataset, Ti, Fe, V, and O. Materials of the same unit cell formula as the transition metal oxide dataset were excluded for the sake of preventing overlap between the two datasets. This means that materials with the same composition as one in the transition metal oxide set, even if they had a unique geometry, were not included in this dataset. This dataset includes 50,000 materials used for training, and 10,000 for validation. A smaller subset of this dataset was created, also excluding Mn and Cr for the purpose of transfer learning with the extended transition metal oxide dataset.

4.2 Transfer Learning Schemes

We investigated three transfer learning schemes. These either chose to freeze weights, not updating them in the new training, or to fine-tune them by training with the new dataset. The schemes used are as follows (see Fig. 7):

1. All weights are fine-tuned (TL1).

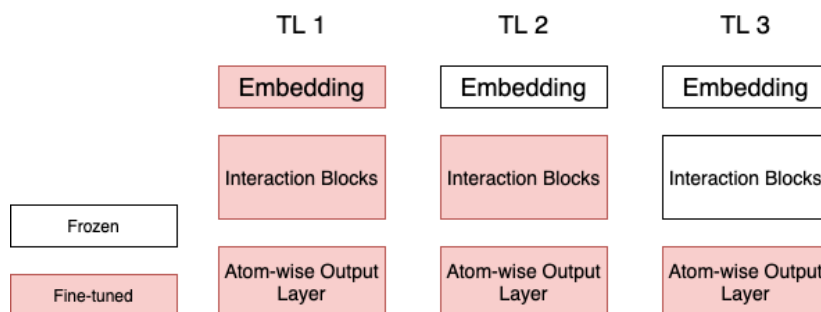


FIG. 7: Visualization of the three transfer learning schemes we implemented. The embedding layer, the interaction blocks, and the output layers are either frozen or fine-tuned.

2. The embedding layer is frozen and the rest of the weights are fine-tuned (TL2).
3. Only the output layers are fine-tuned and all other weights are frozen (TL3).

To obtain a performance benchmark for the target dataset, SchNet is trained with a random initialization on a training set of 400 transition metal oxides. Two different model architectures and regularization schemes were used. Both model architectures used for training on a baseline dataset are based on SchNetPack (Schütt et al., 2019), with and without L2 regularization on all weights. The first one consists of six interaction blocks, 128 length embedding vectors, and 128 filter convolutional layers. Since the transition metal oxide dataset is considerably smaller than the baseline, which consisted of 130,000 materials, a smaller architecture was also trained with and without regularization in an attempt to prevent potential overfitting. This architecture had four interaction blocks, with embeddings of length 30 and convolutional layers with 30 filters. The results of the validation set evaluated by these models are reported in Table 1. The model with the best mean absolute error was the original architecture with an L2 regularization coefficient of 10^{-3} . Out of the validation set, only 47 of the 117 predictions were within chemical accuracy. This best-performing architecture and regularization is used in comparison with transfer learning models.

5. PREDICTION WITH SCHNET TRANSFER LEARNING

Transfer learning methods are compared to direct training methods on the same dataset for varying training set sizes. The target data were split into a training set of 400 and a validation set of

TABLE 1: Validation set mean absolute error in electronvolts of SchNet models trained on the full transition metal oxide training set. Arch 1 is the larger SchNetPack architecture, which consists of six interaction blocks, 128 length embedding vectors, and 128 filter convolutional layers. Arch 2 is the smaller architecture, with four interaction blocks, with embeddings of length 30, and convolutional layers with 30 filters. L2 signifies the inclusion of L2 regularization with a coefficient of 10^{-3} in the loss function during training

N	arch 1	arch 1 + L2	arch 2	arch 2 + L2
400	0.563	0.551	1.175	1.005

117 that remains the same for all evaluations. Training data sizes of 400, 200, and 100 are used. For the 200-length dataset training, the 400 are split into two sets, and a separate neural network is trained with each. The same is done with the 100-length training set, where four neural networks are trained. The mean absolute error on the validation set of each of the same-length dataset networks is then averaged to get the given results. This is done to ensure consistent results, since the dataset is small and all portions may not be entirely representative of each other.

Here, transfer learning methods TL1 and TL2 proved to be the most effective, achieving lower mean absolute error on the validation set than the best direct training model for the 400 and 200 size training sets, and similar error for the 100 size training sets. The third transfer learning scheme performed worse than the direct training. The mean absolute errors are reported in Table 2. A significant difference was seen between the number of validation predictions within chemical accuracy of the DFT value for the successful transfer learning methods compared to direct training, with over 20% more for each training dataset size. These results are visualized in Fig. 8.

Next, the first two transfer learning methods (TL1 and TL2) are compared to direct training on an extended version of the transition metal oxide dataset, which includes two additional transition metal elements, manganese and chromium, and 146 additional data points. On the extended set, training is done similarly except with training set sizes of 500, 250, and 125 and a

TABLE 2: Validation set mean absolute error in electronvolts of each of the transfer learning schemes along with the best-performing direct training method across the different-sized splits of the dataset. The bold numbers are the best result for the row

N	No TL	TL1	TL2	TL3
100	0.939	0.929	0.943	1.191
200	0.762	0.691	0.703	0.847
400	0.551	0.477	0.482	0.686




FIG. 8: Results of the three best transfer learning methods compared to best direct training method without transfer learning with the original dataset. Left: Mean absolute error (MAE) of the validation set evaluation for the models trained on each of the training set sizes. Right: Percent of evaluations within chemical accuracy of the value computed with DFT.

validation set of 163 samples. Like the smaller dataset, the transfer learning methods both get significantly more predictions from the validation set within chemical accuracy of their DFT-calculated value than the directly trained model. These results are reported in Table 3 and Fig. 9. In these two experiments, transfer learning methods TL1 and TL2 perform as good as or better than direct training in mean absolute error. However, even in the two cases with similar mean absolute error, a much larger fraction of the transfer learning predictions are within chemical accuracy than the direct model predictions. The reason for this was explored further, and it was found that in addition to the lower error, the transfer learning predictions also shared a higher proportion of the higher errors than the directly trained model, leading to similar mean error. This is shown for the TL1 model in both the 100 length training sets for the original dataset and the 500 length training set for the extended dataset in Fig. 10. The cumulative error distributions of the TL1 model for the larger and smaller training set size of both the original and extended training data are explored in Fig. 11.

5.1 Computational Cost

As we mentioned in Section 2, DFT calculations can be computationally expensive, with the average calculation from a set of 1300 transition metal oxides taking over 27 minutes per run

TABLE 3: Validation set mean absolute error in electronvolts of the best performing direct training and transfer learning methods applied to the extended metal oxides dataset

N	No TL	TL1	TL2
125	1.355	1.083	1.067
250	0.827	0.753	0.751
500	0.730	0.744	0.730




FIG. 9: Results of the two best transfer learning methods (TL1 and TL2) compared to best direct training method without transfer learning with the extended metal oxides dataset. Left: Mean absolute error (MAE) of the validation set evaluation for the models trained on each of the training set sizes. Right: Percent of model evaluations within chemical accuracy (1 kcal/mol) of the benchmark value computed with DFT. It is seen that transfer learning improves significantly the number of neural net predictions within chemical accuracy, while minimizing MAE.

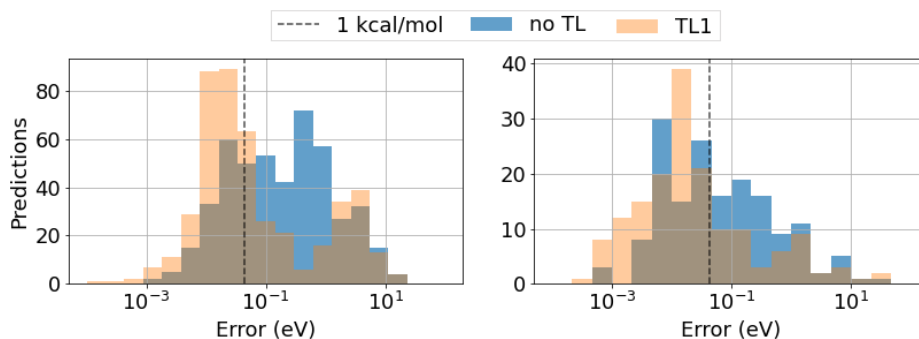


FIG. 10: Error distribution in cases where TL1 and direct training had similar MAE with a significant gap in predictions within chemical accuracy. In each case, TL1 errors make up a larger portion of both the ends of the distribution, with the larger values making a significant impact to the MAE. Left: Predictions from models trained on 100 length training sets of original dataset. Right: Predictions from models trained on 500 length training set of extended dataset.

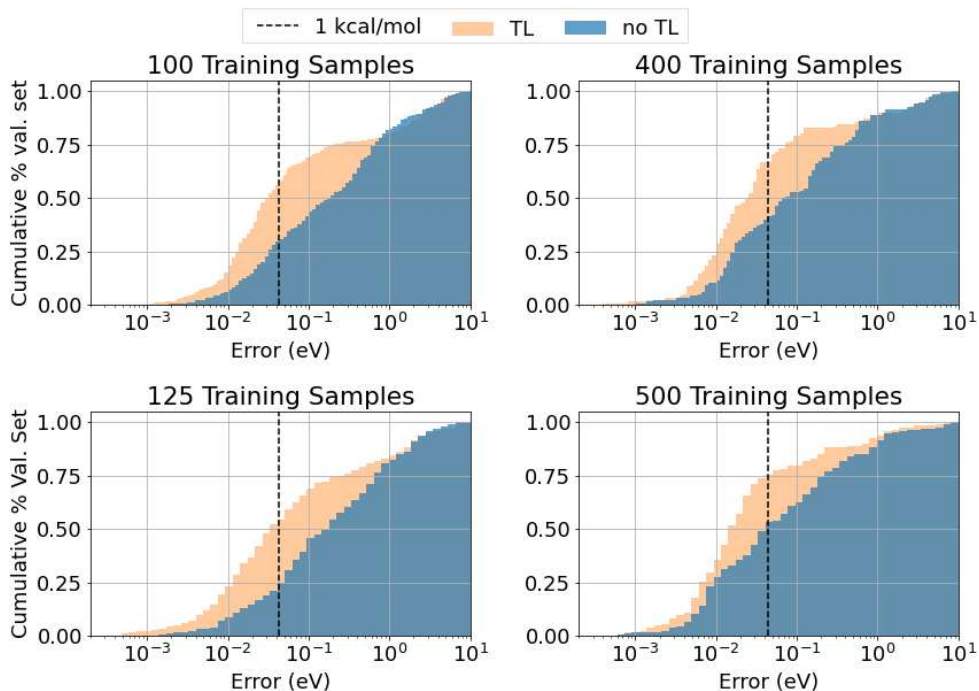


FIG. 11: Top: Cumulative distribution of validation set error for direct training and TL1 methods trained with the 100 and 400 length training sets. In each case, over 20% more of the predictions from transfer learning are within chemical accuracy (1 kcal/mol). Bottom: Cumulative distribution of validation set error for direct training and TL2 methods trained with the 125 and 500 length training sets of the extended data.

(see Fig. 3). Using a neural network in lieu of DFT allows for accelerated predictions after the overhead cost of training and generation of data. With a trained SchNet model, predicting the free energy of a new transition metal oxide is in the order of milliseconds, while running DFT calculations takes similar time as the previous calculations. The overhead cost of training the

neural network can be insignificant for larger scale screening of materials. Table 4 summarizes the training times for direct and transfer learning models. SchNet was implemented using the PyTorch machine learning framework in python (Paszke et al., 2019). Training was done with an NVIDIA Titan RTX GPU. While both direct and transfer training times were similar, we must also consider the training of the source model used to initialize the transfer learning models, which leads to the large discrepancy between the two. For the largest training set, the total training time was 2.2 times that of the mean DFT calculation time for direct training, and 128.2 times for transfer learning. If hundreds or thousands of materials are to be screened, neural networks allow a large savings in computational time, as visualized in Fig. 12.

To give an example of this speedup, DFT was done on a simple, two-atom TiO molecule using the Quantum Espresso DFT code on a 2014 Mac Mini with 1.4 GHz Intel Core i5 processor and 4 GB of RAM. This took 87 seconds. Evaluating the same material with the neural network took 4.6 milliseconds on the same computer. While being a big speedup, this does not capture the more significant speedup seen with materials with more atoms. While DFT calculations for larger transition metal oxides took as long as 20 hours in the benchmark, the longest SchNet evaluation time was 350 milliseconds.

6. CONCLUSIONS

Neural networks have become popular for high-throughput screening of material properties, as they provide significantly faster predictions than *ab initio* simulations based on DFT. However,

TABLE 4: Total computational time in seconds to train direct and transfer learning models per size of training set. Transfer learning models also take into account the training time of their source model, accounting for the large difference between methods

N	No TL	TL1
125	7182	214495
250	5008	210736
500	3626	209751




FIG. 12: Projected computational time for screening properties with only DFT and neural networks. Using 500 samples for training, using either direct training or transfer learning, would save 135 or 78 hours, respectively, in computational time over DFT in predicting properties of just 300 more materials.

in order to obtain accurate predictions neural networks require large amounts of DFT training data, which can be computationally expensive to obtain. To overcome this problem, in this paper we developed new transfer learning algorithms based on SchNet (Schütt et al., 2017b, 2019) to repurpose trained neural network models on other materials using only a small amount of additional training data. We demonstrated that the proposed transfer learning algorithms can improve the mean absolute error of the Gibbs free energy predictions by up to 30% compared to direct training. Furthermore, even in cases where the difference in the mean absolute errors are not significant, transfer learning increased the number of predictions within chemical accuracy by 25% to 30%.

ACKNOWLEDGMENT

This research was supported by the NSF-TRIPODS grant 81389-444168.

REFERENCES

- Behler, J. and Parrinello, M., Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces, *Phys. Rev. Lett.*, vol. **98**, no. 14, p. 146401, 2007.
- Blaha, P., Hofstätter, H., Koch, O., Laskowski, R., and Schwarz, K., Iterative Diagonalization in Augmented Plane Wave Based Methods in Electronic Structure Calculations, *J. Comput. Phys.*, vol. **229**, no. 2, pp. 453–460, 2010.
- Born, M. and Oppenheimer, J.R., On the Quantum Theory of Molecules, *Ann. Physik*, vol. **84**, no. 20, pp. 457–484, 1927.
- Cho, H., Venturi, D., and Karniadakis, G.E., Numerical Methods for High-Dimensional Kinetic Equations, in *Uncertainty Quantification for Kinetic and Hyperbolic Equations*, S. Jin and L. Pareschi, Eds., Berlin: Springer, pp. 93–125, 2017.
- Curtarolo, S., Setyawan, W., Wang, S., Xue, J., Yang, K., Taylor, R.H., Nelson, L.J., Hart, G.L.W., Sanvito, S., Buongiorno-Nardelli, M., Mingo, N., and Levy, O., AFLOWLIB.ORG: A Distributed Materials Properties Repository from High-Throughput *Ab Initio* Calculations, *Comput. Mater. Sci.*, vol. **58**, pp. 227–235, 2012.
- Dektor, A. and Venturi, D., Dynamically Orthogonal Tensor Methods for High-Dimensional Nonlinear PDEs, *J. Comput. Phys.*, vol. **404**, p. 109125, 2020.
- Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning*, Cambridge, MA: MIT Press, 2016.
- Han, J., Lu, J., and Zhou, M., Solving High-Dimensional Eigenvalue Problems Using Deep Neural Networks: A Diffusion Monte Carlo Like Approach, pp. 1–15, 2020. arXiv: 2002.02600
- Hohenberg, P. and Kohn, W., Inhomogeneous Electron Gas, *Phys. Rev.*, vol. **136**, no. 3B, pp. B864–B871, 1964.
- Jain, A., Ong, S.P., Hautier, G., Chen, W., Richards, W.D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., and Persson, K., The Materials Project: A Materials Genome Approach to Accelerating Materials Innovation, *APL Mater.*, vol. **1**, no. 1, p. 011002, 2013.
- Khoromskaja, V. and Khoromskij, B.N., *Tensor Numerical Methods in Quantum Chemistry*, Berlin: De Gruyter, 2018.
- Kingma, D.P. and Ba, J., Adam: A Method for Stochastic Optimization, pp. 1–15, 2017. arXiv: 1412.6980
- Martin, R.M., *Electronic Structure: Basic Theory and Practical Methods*, Cambridge, UK: Cambridge University Press, 2004.

- Pan, S.J. and Yang, Q., A Survey on Transfer Learning, *IEEE Transact. Knowledge Data Eng.*, vol. **22**, no. 10, pp. 1345–1359, 2010.
- Parr, R.G. and Weitao, Y., *Density-Functional Theory of Atoms and Molecules*, Oxford, UK: Oxford University Press, 1994.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S., Pytorch: An Imperative Style, High-Performance Deep Learning Library, in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, Eds., Red Hook, NY: Curran Associates Inc., pp. 8024–8035, 2019.
- Razavian, A.S., Azizpour, H., Sullivan, J., and Carlsson, S., CNN Features Off-the-Shelf: An Astounding Baseline for Recognition, in *The IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, Columbus, Ohio, 2014.
- Reddi, S.J., Kale, S., and Kumar, S., On the Convergence of Adam and Beyond, in *Proc. of the 6th International Conference on Learning Representations (ICLR)*, Vancouver, Canada, 2019.
- Schütt, K.T., Arbabzadah, F., Chmiela, S., Müller, K.R., and Tkatchenko, A., Quantum-Chemical Insights from Deep Tensor Neural Networks, *Nat. Commun.*, vol. **8**, no. 1, p. 13890, 2017a.
- Schütt, K.T., Kindermans, P.J., Felix, H.E.S., Chmiela, S., Tkatchenko, A., and Müller, K.R., SchNet: A Continuous-Filter Convolutional Neural Network for Modeling Quantum Interactions, in *Advances in Neural Information Processing Systems 30*, Red Hook, NY: Curran Associates, Inc., pp. 991–1001, 2017b.
- Schütt, K.T., Kessel, P., Gastegger, M., Nicoli, K.A., Tkatchenko, A., and Müller, K.R., SchNetPack: A Deep Learning Toolbox for Atomistic Systems, *J. Chem. Theor. Comput.*, vol. **15**, no. 1, pp. 448–455, 2019.
- Smith, J., Nebgen, B.T., Zubatyuk, R., Lubbers, N., Devereux, C., Barros, K., Tretiak, S., Isayev, O., and Roitberg, A.E., Approaching Coupled Cluster Accuracy with a General-Purpose Neural Network Potential through Transfer Learning, *Nat. Commun.*, vol. **10**, no. 2903, pp. 1–8, 2019.
- Venturi, D., The Numerical Approximation of Nonlinear Functionals and Functional Differential Equations, *Phys. Rep.*, vol. **732**, pp. 1–102, 2018.
- Venturi, D. and Dektor, A., Spectral Methods for Nonlinear Functionals and Functional Differential Equations, pp. 1–34, 2020. arXiv: 1909.07960
- Wilson, A.C., Roelofs, R., Stern, M., Srebro, N., and Recht, B., The Marginal Value of Adaptive Gradient Methods in Machine Learning, in *Advances in Neural Information Processing Systems 30*, Red Hook, NY: Curran Associates, Inc., pp. 4148–4158, 2017.
- Yamada, H., Liu, C., Wu, S., Koyama, Y., Ju, S., Shiomi, J., Morikawa, and Yoshida, R., Predicting Materials Properties with Little Data Using Shotgun Transfer Learning, *ACS Central Sci.*, vol. **5**, no. 10, pp. 1717–1730, 2019.
- Zinn-Justin, J., *Quantum Field Theory and Critical Phenomena*, 4th Edition, Oxford, UK: Oxford University Press, 2002.