

# Wavelet scattering networks for atomistic systems with extrapolation of material properties

Cite as: J. Chem. Phys. **153**, 084109 (2020); <https://doi.org/10.1063/5.0016020>

Submitted: 31 May 2020 . Accepted: 29 July 2020 . Published Online: 25 August 2020

Paul Sinz , Michael W. Swift , Xavier Brumwell , Jialin Liu , Kwang Jin Kim, Yue Qi , and Matthew Hirn 



View Online



Export Citation



CrossMark

## ARTICLES YOU MAY BE INTERESTED IN

[Optimization of the Reax force field for the lithium–oxygen system using a high fidelity charge model](#)

The Journal of Chemical Physics **153**, 084107 (2020); <https://doi.org/10.1063/5.0014406>

Lock-in Amplifiers  
up to 600 MHz



Watch



# Wavelet scattering networks for atomistic systems with extrapolation of material properties

Cite as: J. Chem. Phys. 153, 084109 (2020); doi: 10.1063/5.0016020

Submitted: 31 May 2020 • Accepted: 29 July 2020 •

Published Online: 25 August 2020



View Online



Export Citation



CrossMark

Paul Sinz,<sup>1</sup>  Michael W. Swift,<sup>2</sup>  Xavier Brumwell,<sup>1</sup>  Jialin Liu,<sup>2</sup>  Kwang Jin Kim,<sup>2</sup> Yue Qi,<sup>2,a)</sup>   
and Matthew Hirn<sup>1,3,4,b)</sup> 

## AFFILIATIONS

<sup>1</sup>Department of Computational Mathematics, Science and Engineering, Michigan State University, East Lansing, Michigan 48824-1226, USA

<sup>2</sup>Department of Chemical Engineering and Materials Science, Michigan State University, East Lansing, Michigan 48824-1226, USA

<sup>3</sup>Department of Mathematics, Michigan State University, East Lansing, Michigan 48824-1226, USA

<sup>4</sup>Center for Quantum Computing, Science and Engineering, Michigan State University, East Lansing, Michigan 48824-1226, USA

**Note:** This paper is part of the JCP Special Topic on Machine Learning Meets Chemical Physics.

<sup>a)</sup>Electronic mail: [yueqi@egr.msu.edu](mailto:yueqi@egr.msu.edu)

<sup>b)</sup>Author to whom correspondence should be addressed: [mhirn@msu.edu](mailto:mhirn@msu.edu)

## ABSTRACT

The dream of machine learning in materials science is for a model to learn the underlying physics of an atomic system, allowing it to move beyond the interpolation of the training set to the prediction of properties that were not present in the original training data. In addition to advances in machine learning architectures and training techniques, achieving this ambitious goal requires a method to convert a 3D atomic system into a feature representation that preserves rotational and translational symmetries, smoothness under small perturbations, and invariance under re-ordering. The atomic orbital wavelet scattering transform preserves these symmetries by construction and has achieved great success as a featurization method for machine learning energy prediction. Both in small molecules and in the bulk amorphous  $\text{Li}_a\text{Si}$  system, machine learning models using wavelet scattering coefficients as features have demonstrated a comparable accuracy to density functional theory at a small fraction of the computational cost. In this work, we test the generalizability of our  $\text{Li}_a\text{Si}$  energy predictor to properties that were not included in the training set, such as elastic constants and migration barriers. We demonstrate that statistical feature selection methods can reduce over-fitting and lead to remarkable accuracy in these extrapolation tasks.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0016020>

## I. INTRODUCTION

Machine learning (ML) is a powerful tool in chemical physics.<sup>1,2</sup> Both kernel-based<sup>3–14</sup> and neural-network-based<sup>15–24</sup> learning algorithms have found success in predicting physical properties such as energies, forces, and potential energy surfaces starting from atomic coordinates. ML models have been used for molecular dynamics (MD),<sup>13,25,26</sup> prediction of free energy surfaces,<sup>27–30</sup> and generation of thermodynamic ensembles<sup>25,31</sup> on systems for which they have been trained. Much as ML models have revolutionized fields such as computer vision,<sup>32</sup> automated content generation,<sup>33</sup>

and natural language processing,<sup>34</sup> an ML model could, in principle, predict the physical properties of broad classes of atomic systems with an accuracy competitive with the best current methods at a small fraction of the computational cost. However, such an ML model has not yet been developed, and many obstacles still remain before general atomistic ML models can compete with the existing quantum chemistry methods. The most fundamental obstacle is perhaps the generalizability problem (also referred to as transferability). Quantum chemical methods such as density functional theory (DFT) are predictive because they work from first principles—physical properties emerge from solutions to equations that describe

underlying physics—allowing these techniques to work on systems that have never been studied before. In contrast, machine learning is at its heart a fitting technique. The model knows nothing about the physical equations, and its predictions are statistical inferences based on the training data. This raises the question: are quantum chemistry models based on machine learning inherently limited to the interpolation of the training data? Or can a machine learning model “learn the underlying physics” and provide new insights beyond the training data? If we want to answer a question with an ML model, will it always be necessary to compute the answer to the same question beforehand in thousands or tens of thousands of similar cases? Or could ML truly teach us something new? The hope that it can is not entirely without foundation. Machine learning models have shown the ability to capture and generalize abstract patterns within training data. For example, a neural net trained to translate languages recently achieved “zero-shot translation”: translation between a pair of languages for which it had no dictionary without going through an intermediate language.<sup>35</sup> This suggests that learning of underlying meaning from diverse examples may be possible.

Of course, the full answer to this question is beyond the scope of the present work. We look at a specific example of the generalizability problem: energy prediction for amorphous  $\text{Li}_\alpha\text{Si}$ . Due to its potential for developing high energy density lithium-ion batteries, this system has recently been studied by a variety of machine learning methods. Onat *et al.* generated an implanted neural network potential for this system.<sup>27</sup> Artrith *et al.* developed a machine-learning potential that enabled ensemble generation and molecular dynamics for  $\text{Li}_\alpha\text{Si}$ .<sup>25</sup> Brumwell *et al.* created a three-dimensional ML model, similar to a convolutional neural network, with a physically motivated filter based on the wavelet scattering transform and were able to achieve chemical accuracy in energy prediction for this class of structures.<sup>37</sup> The wavelet formulation has a number of advantages. The inclusion of wavelet dilations and second-order scattering coefficients makes the wavelet scattering approach inherently multiscale, including length scales on the order of the support of the smallest wavelet up to the length scale of the largest unit cell in the  $\text{Li}_\alpha\text{Si}$  database. This allows the wavelet scattering transform to capture a broad class of physical interactions, making it more general (see Sec. II D for details on the multiscale nature). The fact that the wavelets themselves are based on atomic orbitals may allow them to more naturally capture interactions arising from electronic bonds (though without sacrificing generality since the wavelet frame is overcomplete) and may help with the generalizability of the model. Other methods that sum up atomic energies based on their local environments (e.g., Refs. 7, 25, 27, 38, and 39) can similarly separate length scales by setting the cutoff to the length scale of the unit cell and selecting a variety of radial functions with support along the intermediate length scales down to the distance between nearest-neighbor atoms. These methods provide feature descriptions that separate length scales but must be used in an appropriate training or energy fitting framework in order to nonlinearly couple these scales. In contrast, the wavelet scattering coefficients used here include the nonlinear coupling in the feature set.

In this work, we build on and improve the wavelet-based model,<sup>37,40–42</sup> achieving similar accuracy to Onat *et al.*<sup>27</sup> and Artrith, Urban, and Ceder,<sup>25</sup> but using a simpler architecture with fewer

parameters than the neural network approaches and with fewer features than any of the previous models including Brumwell *et al.*<sup>37</sup> More importantly, we test the model on extrapolation tasks, predicting physical properties that were not present in the training set. We perform detailed analysis on how to balance under-fitting and over-fitting in order to achieve high generalizability. The three tasks are predicting migration barriers based on transition state theory, energies of amorphous systems significantly larger than the training set systems, and elastic properties based on deformations of amorphous  $\text{Li}_\alpha\text{Si}$ . In each of these extrapolation tasks, we find that the model is able to achieve reasonable accuracy, and in some cases does quite well, thus providing evidence of the model’s ability to generalize to new types of systems and tasks.

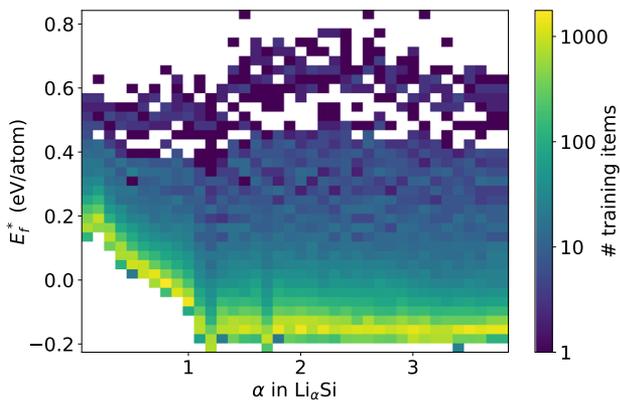
The remainder of this paper is organized as follows: In Sec. II, we present the methods used in this work, including the data generation process and descriptions of the machine learning algorithms. Numerical results are presented and discussed in Sec. III, and Sec. IV contains a short conclusion. Appendixes A–C and supplementary material provide additional details on certain aspects of this paper.

## II. METHODS

Methods consist of data generation for training (Sec. II A) as well as data generation for testing on extrapolation tasks (Sec. II B). Algorithms used for training the machine learned models are discussed in Sec. II C, whereas the wavelet scattering representation of an atomic state is described in Sec. II D. Appendix B explains how to compute such features efficiently.

### A. Training data generation

The training, validation, and interpolation testing data for the machine-learned model consist of amorphous  $\text{Li}_\alpha\text{Si}$  structures labeled by formation energies calculated using Density Functional Theory (DFT). These structures are in cubic boxes under periodic boundary conditions containing from 55 to 100 atoms, with lithium-to-silicon ratio  $\alpha$  ranging from 0.1 to 3.75. The initial disordered structures are generated by evolving random structures under ReaxFF<sup>43</sup> molecular dynamics (MD) at 2500 K for 10 ps, and ten different disordered structures are randomly picked from the MD trajectory for each of the 37 chosen concentrations. The accuracy of the force field used to obtain the initial amorphous structure is not important due to the following DFT calculations. In particular, each structure is fully relaxed at constant volume using DFT. The structures and formation energies along the relaxation paths make up the amorphous dataset used in this work, which contains a total of 90 252 structures. A histogram of the quantity of these structures by energy and concentration is shown in Fig. 1. We note that the structures are heavily concentrated near the endpoint of the relaxation, so we expect the resulting model to do better on near-equilibrium amorphous structures. This is desirable because the low-energy structures are more likely to arise in realistic simulations. We also calculate voltages vs  $\text{Li}/\text{Li}^{+44–46}$  and radial distribution functions for Si–Si, Li–Si, and Li–Li pairs and find good agreement with similar datasets from the literature<sup>25,47</sup> (see the supplementary material). This confirms that our amorphous structures are physically realistic.



**FIG. 1.** Histogram of training set energies vs concentration  $\alpha$  in  $\text{Li}_\alpha\text{Si}$ . Color indicates the number of training items in each bin on a logarithmic scale.

Formation energies and relaxations were performed in the Vienna *Ab initio* Simulation Package (VASP)<sup>48</sup> using the Projector-Augmented Wave method<sup>49</sup> and the PBE exchange–correlation functional<sup>50</sup> with a plane-wave energy cutoff of 500 eV. The Brillouin zone was sampled using the Gamma point only during relaxation. After relaxation, the energies along each relaxation path were corrected for  $k$ -point sampling errors by calculating the energy of each fully relaxed structure using a  $3 \times 3 \times 3$  Gamma-centered grid and applying the resulting constant shift to the rest of the structures in the relaxation path. The mean absolute  $k$ -point sampling correction was 27 meV/atom. The total formation energy of a structure with  $N_{\text{Li}}$  lithium atoms and  $N_{\text{Si}}$  silicon atoms is defined based on DFT total energies,

$$E_f(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}}) = E_{\text{tot}}(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}}) - N_{\text{Li}}E(\text{Li}) - N_{\text{Si}}E(\text{Si}),$$

where  $E_{\text{tot}}(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}})$  is the total energy of the system and  $E(\text{Li})$  and  $E(\text{Si})$  are the DFT total energy per atom of elemental lithium and silicon, respectively. The structure  $\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}}$  has reduced formula  $\text{Li}_\alpha\text{Si}$  with  $\alpha = N_{\text{Li}}/N_{\text{Si}}$  and per-atom formation energy,

$$E_f^*(\text{Li}_\alpha\text{Si}) = E_f(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}}) / (N_{\text{Li}} + N_{\text{Si}}). \quad (1)$$

The per-atom formation energy is the quantity of interest for machine learning. Notice, though, it includes the terms  $N_{\text{Li}}E(\text{Li})$  and  $N_{\text{Si}}E(\text{Si})$ , which require no additional quantum mechanical calculations beyond the one-time cost of computing  $E(\text{Li})$  and  $E(\text{Si})$ . The difficulty is in computing  $E_{\text{tot}}(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}})$ , which requires a costly DFT calculation for each new state. When fitting our machine learned models, we regress the per-atom total energy, defined as

$$E_{\text{tot}}^*(\text{Li}_\alpha\text{Si}) = E_{\text{tot}}(\text{Li}_{N_{\text{Li}}}\text{Si}_{N_{\text{Si}}}) / (N_{\text{Li}} + N_{\text{Si}})$$

or

$$E_{\text{tot}}^*(\text{Li}_\alpha\text{Si}) = E_f^*(\text{Li}_\alpha\text{Si}) + \frac{\alpha}{1 + \alpha}E(\text{Li}) + \frac{1}{1 + \alpha}E(\text{Si}).$$

Even though it is simple to convert total energies into per-atom total energies, we regress the latter since per atom energies remove

the effect of varying unit cell sizes and the number of atoms per unit cell on the total energy. Since we use the squared loss as our measure of error when training, regressing total energies would bias the models toward systems containing larger numbers of atoms since the total energy scales with the number of atoms.

## B. Data generation for extrapolation tests

In order to test the machine learning model's generalizability to extrapolation tasks, additional DFT data are required to compare with the results of the machine learning model. We test three different extrapolation tasks: prediction of migration barriers, energy prediction for systems with larger unit cells, and prediction of elastic properties.

Diffusion barriers cannot be defined uniquely in amorphous structures due to the lack of order. Rather, paths that move an atom from one favorable coordination environment to another through a relatively unfavorable environment are abundant. An endpoint for such a pathway was found by locating void spaces in the amorphous structure through Voronoi analysis and inserting a test lithium atom at each void to find the most energetically favorable position. Nearby lithium atoms to this void were subsequently identified, and the minimum-energy path (MEP) for each lithium to travel to the void was calculated using the nudged elastic band (NEB) method.<sup>51</sup> The NEB images along 6 calculated MEPs (for a total of 50 image structures) were used as testing data for this extrapolation task. The primary quantity of interest is the migration barrier, i.e., the difference between the lowest-energy and highest-energy points along the MEP.

Large structure testing data were generated by two methods: independently relaxing larger AIMD-generated structures (the “from-scratch” method) or tiling structures from the dataset, randomly perturbing all atomic positions by 0.1 Å, and performing a single-point calculation (the “tiled” approach). The testing data consist of 37 from-scratch structures,  $40 \times 2 \times 2 \times 2$  tiled structures, and  $108 \times 2 \times 1 \times 1$  tiled structures.

Finally, elastic property testing data were generated by applying hydrostatic strain from  $-9\%$  to  $+9\%$  on fully relaxed structures at each concentration. The bulk modulus  $K$  is calculated by fitting data near the minimum to the following equation:<sup>52</sup>

$$K = V_0 \frac{\partial^2 E}{\partial V^2},$$

where  $V$  is the volume,  $V_0$  is the equilibrium volume, and  $E$  is the energy. In total, a bulk modulus value is calculated at each of the 37 concentrations, based on a total of 333 structures under hydrostatic strain.

## C. Linear regression and model fitting

The purpose of using machine learning is to reduce the computational burden relative to quantum mechanical calculations while maintaining accurate predictions. To accomplish this task, we need to derive features from our data and recombine them in some meaningful way. We use a linear model over a predefined set of features, which consist of nonlinear, multiscale maps of the original atomic state. Using a linear model over a universal feature set allows us to leverage several well-studied techniques in regression,

regularization, and statistical learning theory that increases the accuracy, stability, and generalizability of the model.

Let  $x = \{(Z_k, R_k) \in \mathbb{N} \times \mathbb{R}^3\}_{k=1}^{N_x}$  denote the list of atoms in the unit cell of a  $\text{Li}_\alpha\text{Si}$  system. The value  $Z_k \in \mathbb{N}$  denotes the protonic charge of the atom, i.e.,  $Z_k = 3$  for lithium and  $Z_k = 14$  for silicon, and  $R_k \in \mathbb{R}^3$  denotes the position of the atom in the simulation cell. The quantity  $N_x = N_{\text{Li}} + N_{\text{Si}}$  is the total number of atoms in the unit cell. Let  $\Phi(x) \in \mathbb{R}^d$  be a  $d$ -dimensional representation of the state  $x$ , which is described in detail in Sec. II D. A linear regression with weights  $w = (w_\gamma)_{\gamma=0}^d \in \mathbb{R}^{d+1}$  of the per-atom total energy  $E_{\text{tot}}^*(x)$  over the representation  $\Phi(x) = (\phi_\gamma(x))_{\gamma=1}^d$  computes

$$\tilde{E}_{\text{tot}}^*(x; w) = \tilde{E}(x; w) = w_0 + \sum_{\gamma=1}^d w_\gamma \phi_\gamma(x), \quad (2)$$

where  $w_0$  is a bias term and the coordinates  $\phi_\gamma(x)$  of  $\Phi(x)$  are weighted with the scalars  $w_\gamma$ . We regularize the regression by selecting a parsimonious model via a sparsity constraint on the weights  $w$ ,

$$\|w\|_0 = \#\{w_\gamma \neq 0 : 0 \leq \gamma \leq d\} \leq M,$$

for some hyper-parameter  $M$  that determines the number of nonzero weights.

The weights  $(w_\gamma)_{\gamma=0}^d$  and the hyper-parameter  $M$  are solved for using the DFT generated training data. Let  $\mathcal{X}_t = \{(x_i, E(x_i))\}_{i=1}^{n_t}$  denote a training set consisting of  $\text{Li}_\alpha\text{Si}$  states  $x_i$  and their DFT generated per-atom total energies  $E(x_i) = E_{\text{tot}}^*(x_i)$ ; denote by  $\mathcal{X}_v = \{(x'_i, E(x'_i))\}_{i=1}^{n_v}$  another such set, also consisting of  $\text{Li}_\alpha\text{Si}$  states and their associated per-atom total energies, non-overlapping with the training set, which we use as the validation set. For each  $M$  up to some maximum value,  $1 \leq M \leq M_{\text{max}}$ , we compute weights  $w^M = (w_\gamma^M)_{\gamma=0}^d$  by solving the following:

$$w^M = \arg \inf_{w \in \mathbb{R}^{d+1}} [L(w, \mathcal{X}_t) : \|w\|_0 \leq M], \quad (3)$$

where  $L(w, \mathcal{X}_t)$  is the mean squared loss function with respect to the training set  $\mathcal{X}_t$ ,

$$L(w, \mathcal{X}_t) = \frac{1}{n_t} \sum_{i=1}^{n_t} |E(x_i) - \tilde{E}(x_i; w)|^2.$$

As  $M$  increases, the model  $\tilde{E}(x; w^M)$  becomes more complex as it has more non-zero weights. This increasing complexity is reflected by the fact that the loss function  $M \mapsto L(w^M, \mathcal{X}_t)$  is a decreasing function of  $M$ . That is, the training error decreases as  $M$  increases (see also the red curves in Fig. 5).

However, it is well known in machine learning and statistical learning theory that more complex models do not necessarily generalize better. The optimal regularization, here controlled by the sparsity parameter  $M$ , is determined via cross-validation using the validation set  $\mathcal{X}_v$ . That is, for each  $1 \leq M \leq M_{\text{max}}$ , we compute the loss of the model  $\tilde{E}(x; w^M)$  on the validation set and select the  $M^*$  that minimizes the validation error,

$$M^* = \arg \inf_{1 \leq M \leq M_{\text{max}}} L(w^M, \mathcal{X}_v). \quad (4)$$

The model that is used on the test data is  $\tilde{E}(x; w^{M^*})$ . In general,  $M^* \neq M_{\text{max}}$ , since unlike the training error  $M \mapsto L(w^M, \mathcal{X}_t)$ , the validation error  $M \mapsto L(w^M, \mathcal{X}_v)$  is not monotonically decreasing but rather generically decreases up to  $M^*$  and then increases after  $M^*$  (see the green dashed curves in Fig. 5). We remark that the value  $M^*$  is the best estimate of the optimal model for testing on states similar to those in the validation set; in other words, it balances the model between under-fitting and over-fitting. However, states for the extrapolation tasks (Sec. II B) are not necessarily similar to the states in the validation set. Models that extrapolate must be formulated in such a way that when trained, the cross-validation procedure selects a complexity  $M^*$  that captures underlying physical phenomena while ignoring non-physical patterns in the training and validation data. We achieve such a result by leveraging the universal wavelet scattering features (see Sec. II D) and by careful partitioning of the training and validation sets, which are described in more detail in Sec. III A.

Computationally, solving (3) is NP-Hard. We thus solve a relaxed problem that obtains the weights in a greedy fashion using the same orthogonal least squares (OLS) approach described by Hirn, Mallat, and Poilvert.<sup>40</sup> While the resulting weights do not in general solve (3) due to the relaxation, the OLS approach is optimal among greedy approaches since it reduces the mean square error by the maximum amount with each greedy step. Otherwise, using a greedy approach has two benefits. First, it is significantly more efficient; solving (3) requires  $O(\binom{d}{M})$  floating point operations, whereas the greedy approach requires  $O(dM)$  floating point operations. Second, it is an iterative process that can be solved using a QR factorization, which, in this case, means that after  $M$  nonzero weights are selected, the computation for  $M + 1$  nonzero weights requires solving only for the one additional weight. This lets us efficiently construct an array of models for  $1 \leq M \leq M_{\text{max}}$ , which, in turn, enables an efficient solution to the cross-validation problem given in (4).

Finally, we augment the learning process by leveraging empirical bootstrapping and feature bagging. Given an initial database of  $n$  states and their energies (that does not include the withheld testing set), the empirical bootstrap algorithm samples the database with replacement to obtain the training set. Those states not selected for the training set are placed in the validation set. This approach allows us to construct many different models from one database, which are then averaged. The resulting averaged model, which is still a linear model over the representation  $\Phi(x)$ , is superior to any one individually fitted model since the averaging reduces random fluctuations in the fitting process that result from spurious patterns in a single training set. In order for this averaging process to have a maximum effect, the weights of the individual models must be as uncorrelated as possible. Feature bagging, which is a prominent component of random forests, decorrelates the models by restricting the greedy selection at each greedy step. In particular, at each greedy step in the OLS algorithm, approximately  $\sqrt{d+1}$  features are sampled without replacement from among the full set of  $d$  features in  $\Phi(x)$  plus the bias term, minus the features that have already been selected up to that point. The OLS algorithm at each step must then select from among the sampled features, which due to the randomness in the feature sampling, results in models that are significantly less correlated. Indeed, in our own numerical experiments, the most

significant features selected with empirical bootstrapping, but without feature bagging, are very often identical. While restricting the number of possible features at each greedy step means that each model has larger error on the training set, the aggregated average model improves on the test set.<sup>53</sup>

#### D. Atomic orbital wavelet scattering

We now describe how we construct the feature vector  $\Phi(x)$ . Since our regression  $\hat{E}(x; w) = w_0 + \langle w, \Phi(x) \rangle$  is a linear model over  $\Phi(x)$ , the representation  $\Phi(x)$  should have the same properties as  $E(x)$ . In particular, as has been noted by several machine learning papers for many-particle physics,<sup>7</sup>  $E(x)$  is invariant to translations, rotations, and reflections (i.e., isometries) of the atomic coordinates  $\{R_k\}_{k=1}^{N_x}$ , and therefore,  $\Phi(x)$  should also be invariant to isometries. Additionally,  $E(x)$  is independent of the atom indexation in the list  $x$ , and thus,  $\Phi(x)$  must be invariant to index permutations. Like the formation energy,  $\Phi(x)$  should be a continuous function of the atomic coordinates  $\{R_k\}_{k=1}^{N_x}$ , which is particularly important since our data consists of structural relaxation paths. Furthermore, since we are fitting total energy per atom, the feature values should be independent of the number of atoms in the unit cell. Finally, the amorphous  $\text{Li}_\alpha\text{Si}$  systems are periodic, and thus, the features  $\Phi(x)$  must be invariant to equivalent representations of  $x$  with respect to the periodicity of the state.

In addition to those basic physical properties, electronic interactions encoded by the molecular Hamiltonian are inherently multiscale in nature and thus range over a number of length scales. The resulting total energy of the system is a nonlinear function of these length scales. We thus seek a representation  $\Phi(x)$  that on the one hand can separate the length scales of the system, while on the other hand can recombine them in a nonlinear fashion. Our approach is to use the atomic orbital wavelet scattering transform of Brumwell *et al.*,<sup>37</sup> which itself is an adaptation of the three-dimensional solid harmonic wavelet scattering transform proposed by Eickenberg *et al.*<sup>41,42</sup> for molecules. We review its construction in this section, emphasizing certain nuances specific to atomic states with periodic unit cell and more specifically to  $\text{Li}_\alpha\text{Si}$  systems.

The wavelet scattering transform is based upon the wavelet transform of three-dimensional signals. We identify the state  $x$  with such a signal, which will encode permutation invariance into all representations derived from this signal. Let  $\mathcal{Q}_x \subset \mathbb{R}^3$  be the unit cell of the state  $x$ , which in the case of all systems we consider is a cube. We encode the state  $x$  as a superposition of Dirac delta functions,

$$\rho_x^f(u) = \sum_{k=1}^{N_x} f(Z_k) \delta(u - R_k), \quad u \in \mathcal{Q}_x.$$

We use the notation  $\rho_x^f$  because one can think of it as a type of nuclear density for the state  $x$ , but in which we allow some additional flexibility. In particular, the function  $f: \mathbb{N} \rightarrow \mathbb{R}$  encodes a weight that responds based on the type of atom. We use five different functions  $f$ , which can be thought of as channels of the state  $x$ , similar to how a color image has red, green, and blue channels. The five channels we use are lithium, silicon, valence, ionic, and kinetic. The lithium and silicon channels partition the state  $x$  along atom species, whereas the valence and core channels separate the state  $x$  according to electron type. The kinetic channel, inspired by the

Thomas–Fermi–Dirac–von Weizsacker model in quantum chemistry, encodes a different scaling in the number of electrons than the other four channels. The precise definitions of these channels are given in Appendix A.

A simple translation and rotation invariant representation of  $x$  is obtained by summing over  $f_x = (f(Z_k))_{k=1}^{N_x}$  for each channel  $f$ ,

$$\phi_\gamma(x) = N_x^{-1} \|f_x\|_q^q = \frac{1}{N_x} \sum_{k=1}^{N_x} |f(Z_k)|^q, \quad \gamma = (f, q). \quad (5)$$

We compute four different types of summations by taking  $q$ th powers of  $f(Z_k)$  for

$$q \in \{1, 4/3, 5/3, 2\}. \quad (6)$$

These powers are also inspired by the Thomas–Fermi–Dirac–von Weizsacker model and will be discussed more later in this section. By dividing by  $N_x$ , the features  $\phi_\gamma(x)$  are also invariant to system size. Finally, since they are constant with respect to the atomic coordinates, they are trivially continuous functions of them.

The zero order features (5) satisfy all the required invariance properties, but they remove all geometric information from the state  $x$  and are constant for a given lithium–silicon ratio  $\alpha$ . We compute complimentary features that separate the length scales of  $x$  and encode multiscale geometric invariants. These features will be derived from a three-dimensional wavelet transform of  $\rho_x^f$ , which gives a multiscale representation of the signal. Following Ref. 37, we define a family of atomic orbital wavelets  $\psi_{n,\ell}^m: \mathbb{R}^3 \rightarrow \mathbb{C}$ ,

$$\psi_{n,\ell}^m(u) = Q_{n,\ell} (|u|) Y_\ell^m(u/|u|), \\ n \geq 1, 0 \leq \ell < n, |m| \leq \ell,$$

where  $Y_\ell^m$  is the usual spherical harmonic function and  $Q_{n,\ell}$  is a radial function defined as

$$Q_{n,\ell}(r) = C_{n,\ell} r^\ell L_{n-\ell-1}^{\ell+1/2} \left( \frac{r^2}{2\beta^2} \right) e^{-r^2/2\beta^2}, \quad r \geq 0. \quad (7)$$

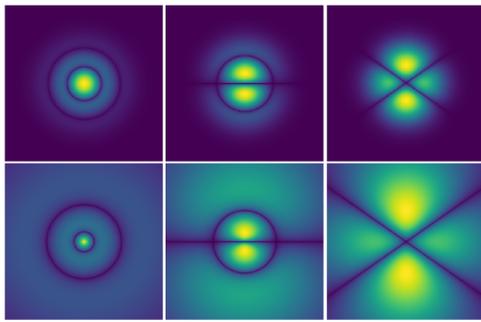
Here,  $C_{n,\ell}$  is a normalizing constant and the functions  $L_k^v$  are the associated Laguerre polynomials. We refer to the family of wavelets  $\psi_{n,\ell}^m$  as atomic orbital wavelets since they mimic the shape of the hydrogen atomic orbitals. Indeed,  $(n, \ell) = (1, 0)$ ,  $(2, 0)$ , and  $(2, 1)$  corresponds to the 1s, 2s, and 2p orbitals, respectively, with similar correspondences for larger values of  $n$  (see Fig. 2). While the hydrogen atomic orbitals have exponential scaling, here, we use a Gaussian function, which mimics the well-known Gaussian type orbitals from the quantum chemistry literature.

We use these wavefunctions as wavelets, though, in which the wavelet transform dilates the wavelet at different dyadic scales  $2^j$ ,

$$\psi_{j,n,\ell}^m(u) = 2^{-3j} \psi_{n,\ell}^m(2^{-j}u), \quad 0 \leq j < J,$$

which increases the size of the wavelet. Let  $s_x$  be the side length of the cubic unit cell  $\mathcal{Q}_x$ . Unlike the molecular systems studied by Eickenberg *et al.*,<sup>41,42</sup> here,  $x$  is a periodic system and so we compute a periodic wavelet transform of the density  $\rho_x^f$  using a circular convolution  $\otimes$  defined as

$$\rho_x^f \otimes \psi_{j,n,\ell}^m(u) = \sum_{p \in \mathbb{Z}} \rho_x^f * \psi_{j,n,\ell}^m(u - ps_x). \quad (8)$$

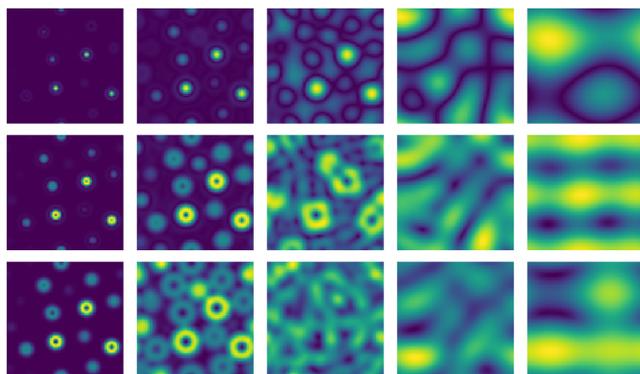


**FIG. 2.** Top row left to right: Density plot cross sections in the  $xz$ -plane of atomic orbital wavelets for  $(n, \ell, m) = (3, 0, 0)$ ,  $(3, 1, 0)$ , and  $(3, 2, 0)$ . Bottom row: Corresponding plots for the hydrogen atom orbitals (3s, 3p, and 3d orbitals, respectively). All images are rescaled for visualization. Note that the exponential radial decay  $e^{-r}$  of the hydrogen orbitals is replaced by a Gaussian decay  $e^{-r^2}$  in the wavelets accounting for the greater localization of the wavelets.

The operation  $*$  is the usual convolution over  $\mathbb{R}^3$ , which for the nuclear density-type function  $\rho_x^f$  yields

$$\rho_x^f * \psi_{j,n,\ell}^m(u) = \sum_{k=1}^{N_x} f(Z_k) \psi_{j,n,\ell}^m(u - R_k). \quad (9)$$

Examining (8) and (9), we have the following interpretations. The standard convolution  $\rho_x^f * \psi_{j,n,\ell}^m$  emits the wavelet  $\psi_{j,n,\ell}^m$  from the location of each atom in the unit cell of  $x$ , with a strength given by  $f(Z_k)$ . The interference patterns encoded by these emissions encode geometric information of the state of the system at different scales  $2^j$ , which we shall aggregate to form multiscale, invariant features. The circular convolution  $\rho_x^f \otimes \psi_{j,n,\ell}^m$  wraps the wavelets periodically in the unit cell  $Q_x$ , thus giving us a periodic function that respects the periodicity of the system. The parameter  $\beta$  in (7) in the definition of  $Q_{n,\ell}$ , which encodes the smallest wavelet scale, is selected



**FIG. 3.** Cross sections of the first order nonlinear, equivariant maps  $\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)(u)$ . The  $\log_2$  scales  $j = 0, 1, 2, 3, 4$  increase from left to right, respectively, and the angular quantum number  $\ell = 0, 1, 2$  from top to bottom, respectively, with  $n = 3$ . The maps extract multiscale geometric information on the arrangement of the atoms in the state  $x$ .

so that  $\psi_{n,\ell}^m(u - R_k)$  and  $\psi_{n,\ell}^m(u - R_l)$  interfere only if  $|R_k - R_l|$  is small, i.e., if the atoms located at  $R_k$  and  $R_l$  are neighboring atoms. The maximum scale  $2^{j-1}$  is selected so that the size of the wavelet  $\psi_{j-1,n,\ell}^m$  is on the order of the maximum side length  $s_x$  of the unit cells  $Q_x$  across all training states, thus enabling the corresponding wavelet coefficients  $\rho_x^f \otimes \psi_{j-1,n,\ell}^m$  to encode macroscopic patterns in the arrangement of the atoms in  $x$ . These choices allow us to capture short-range interactions in features derived from wavelet filters with small  $j$ , while wavelet filters with large  $j$  capture interactions across a larger span of the system.

Convolution operators are, in general, translation equivariant, but not rotation equivariant. However, the atomic orbital wavelet filters are designed to admit a rotationally equivariant representation by combining information across the magnetic quantum number  $m$  through the following nonlinear transform  $\sigma$ :

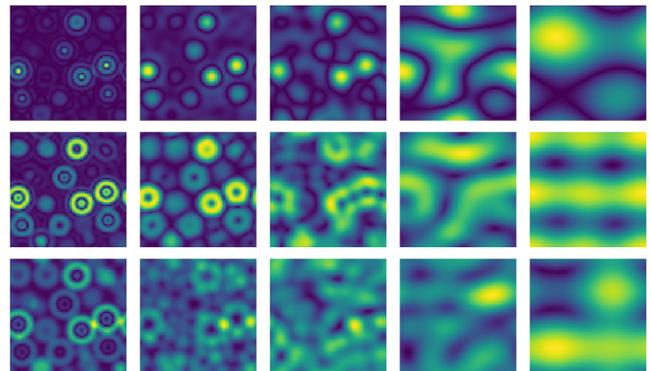
$$\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)(u) = \left( \sum_{m=-\ell}^{\ell} |\rho_x^f \otimes \psi_{j,n,\ell}^m(u)|^2 \right)^{1/2}.$$

The collection of maps  $\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)$  constitutes a multiscale, isometry equivariant representation of the state  $x$  (see Fig. 3 for plots of these maps). We note that the wavelet filters with  $\ell = 0$  are radially symmetric, and therefore, their convolutions are equivariant to rotations with or without the nonlinear transform  $\sigma$ .

Equivariant representations yield invariant representations via integral operators that integrate over the space variable  $u$ . We compute  $L^q(Q_x)$  norms, for the same four  $q$  values in (6), of the maps  $\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)$ ,

$$\|\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)\|_q^q = \int_{Q_x} |\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m)(u)|^q du. \quad (10)$$

The selection of powers  $q$  is motivated by the Thomas–Fermi–Dirac–von Weizsacker model in quantum chemistry in which the 4/3 scaling is used to approximate the exchange energy, the 5/3 scaling is used to approximate the kinetic energy, and the power of 2 encodes an additional part of the kinetic energy and pairwise



**FIG. 4.** Cross sections of the second order nonlinear, equivariant maps  $\sigma(\sigma(\rho_x^f \otimes \psi_{j_1,n_1,\ell_1}^m) \otimes \psi_{j_2,n_2,\ell_2}^m)(u)$  for  $(j_1, n_1, \ell_1) = (1, 3, 1)$ , which is the second from the top and second from the left in Fig. 3. The  $\log_2$  scales  $j_2$  and  $\ell_2$  vary the same as in Fig. 3, and  $n_2 = 3$ . Note that many of the multiscale geometric patterns are distinct from those in Fig. 3.

Coulombic interactions (see also Ref. 40). The power  $q = 1$  is also used since these integrals scale linearly with  $\sum_k f(Z_k)$ .

We normalize the norms (10) to be invariant to system size, which defines first order wavelet scattering features,

$$\phi_\gamma(x) = N_x^{-1} \|\sigma(\rho_x^f \otimes \psi_{j,n,\ell})\|_q^q, \quad \gamma = (f, q, j, n, \ell).$$

In numerical experiments reported on in Sec. III, there are five channels  $f$ , five scales  $j$  (i.e.,  $J = 5$ ),  $n = 3$  (that is, a single  $n$  is used),  $0 \leq \ell < n = 3$ , and there are the four  $q$  values specified in (6), which yields 300 first order features. The smallest wavelet scale (at  $j = 0$ ) is set with the parameter  $\beta$  in (7) to be large enough to ensure the nontrivial overlap between nearest neighbor atoms in the database, which corresponds to a length scale of 0.9 Å. The largest length scale (at  $j = J - 1 = 4$ ) is chosen so that the wavelet supports envelop the largest unit cell in the database, which has side length 11.8 Å. These first order wavelet scattering features encode isometry and size invariant descriptions of the state  $x$  across multiple length scales  $2^j$  for  $0 \leq j < J$ . Furthermore, since the atomic orbital wavelets  $\psi_{n,\ell}^m$  are continuous functions, the resulting maps  $\sigma(\rho_x^f \otimes \psi_{j,n,\ell})$  are continuous functions of the atomic coordinates, which means that their integrals are as well. The use of circular convolution ensures that the maps are invariant to the representation of  $x$  with respect to its periodicity.

First order wavelet scattering features are complemented by second order wavelet scattering features that incorporate multiple length scales of  $x$  into a single feature. They are computed by iterating the nonlinear wavelet transform, which couples the scales  $2^{j_1}$  and  $2^{j_2}$ ,

$$\sigma(\sigma(\rho_x^f \otimes \psi_{j_1, n_1, \ell_1}) \otimes \psi_{j_2, n_2, \ell_2})(u) = \left( \sum_{m=-\ell_2}^{\ell_2} |\sigma(\rho_x^f \otimes \psi_{j_1, n_1, \ell_1}) \otimes \psi_{j_2, n_2, \ell_2}^m(u)|^2 \right)^{1/2}. \quad (11)$$

The second order maps (11), which resemble the architecture of a convolutional neural network as well as aspects of tensor field networks,<sup>24</sup> are equivariant with respect to translations and rotations, and extract coupled geometric information at the scales  $2^{j_1}$  and  $2^{j_2}$  from the state  $x$ . Figure 4 plots the examples of these maps, which are noticeably different than their first order counterparts.

Second order invariant wavelet scattering features are computed analogously to the first order features, by taking normalized  $L^q(Q_x)$  norms of the equivariant maps,

$$\begin{aligned} \phi_\gamma(x) &= N_x^{-1} \|\sigma(\sigma(\rho_x^f \otimes \psi_{j_1, n_1, \ell_1}) \otimes \psi_{j_2, n_2, \ell_2})\|_q^q \\ &= \frac{1}{N_x} \int_{Q_x} |\sigma(\sigma(\rho_x^f \otimes \psi_{j_1, n_1, \ell_1}) \otimes \psi_{j_2, n_2, \ell_2})(u)|^q du \\ \gamma &= (f, q, j_1, n_1, \ell_1, j_2, n_2, \ell_2). \end{aligned}$$

Using the same parameters as the first order features, plus setting  $n_2 = 3$  and  $0 \leq \ell_2 < n_2 = 3$  and  $\max(0, j_1 - 1) \leq j_2 < J = 5$ , we see that there are 3420 second order wavelet scattering features. They thus greatly expand the representation of the state  $x$  and satisfy all the required invariance properties.

We collect the zero, first, and second order invariant wavelet scattering features into a single feature representation  $\Phi(x) \in \mathbb{R}^{3740}$ .

Let  $\lambda = (j, n, \ell)$  denote the triplet of wavelet parameters. With this feature representation, our energy model (2) can be written as

$$\begin{aligned} \tilde{E}_{\text{tot}}^*(x; w) &= \tilde{E}(x; w) = w_0 + \frac{1}{N_x} \sum_{f,q} w_{f,q} \|f_x\|_q^q \\ &+ \frac{1}{N_x} \sum_{f,q,\lambda} w_{f,q,\lambda} \|\sigma(\rho_x^f \otimes \psi_\lambda)\|_q^q \\ &+ \frac{1}{N_x} \sum_{f,q,\lambda_1,\lambda_2} w_{f,q,\lambda_1,\lambda_2} \|\sigma(\sigma(\rho_x^f \otimes \psi_{\lambda_1}) \otimes \psi_{\lambda_2})\|_q^q, \quad (12) \end{aligned}$$

with the weights solved for using the algorithms described in Sec. II C. We remark that both the features themselves, and the number of features, constitute a vast simplification over the model presented by Brumwell *et al.*<sup>37</sup>

### III. NUMERICAL RESULTS

We report empirical errors for formation energy prediction and related tasks. Section III A presents interpolation type test errors on predicting relaxation path formation energies for amorphous Li<sub>α</sub>Si of a similar nature to those in the training set. Sections III B–III D report extrapolation task errors, focusing on diffusion barrier estimation, formation energies of amorphous states with larger unit cells, and bulk modulus prediction, respectively. Appendix C compares these results with results obtained by varying the training procedures and model formulation presented in Sec. II. In particular, it considers a non-randomized training procedure that does not utilize bootstrapping and feature bagging (originally described in Sec. II C), as well as an energy model that utilizes only zero and first order wavelet scattering features (defined in Sec. II D). In each case, we see the advantage of the full algorithm.

#### A. Training and testing on amorphous dataset

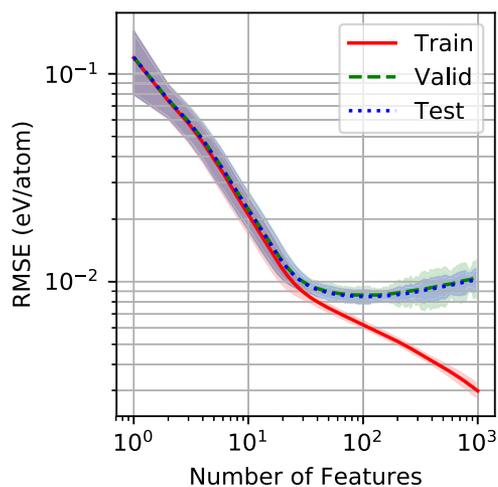
Recall from Sec. II A and Fig. 1, we have a training database of 90 252 amorphous Li<sub>α</sub>Si structures with DFT computed energies spread across 37 concentrations  $\alpha$ , ranging from 0.1 to 3.75. These 90 252 structures correspond to 370 relaxation paths of an initial set of 370 high energy states, with 10 relaxation paths per concentration. Unlike many standard machine learning approaches that obtain a training set by uniformly sampling data points from the initial database, here we uniformly randomly sample relaxation paths. Using fivefold cross-validation, we randomly partition the relaxation paths into five sets of 74 relaxation paths with two paths per concentration in each of the sets of 74. We place four of these sets, 296 relaxation paths total, in the training/validation set, and one set of 74 paths in the test set. We rotate through using each set separately as a test set, meaning that we carry out all numerical experiments five times, each time with a different training/validation and test set split.

We select the training/validation/testing sets according to relaxation data paths, and not structures, because it leads to more physically realistic training and testing scenarios. In particular, it is reasonable to assume that whole relaxation paths, computed via DFT, would be used to train a model, which is then used to compute relaxation paths of new high energy states. Here, the validation and

testing sets are simpler in that we require the model to predict all formation energies along a new relaxation path in which the structures along the path are given. Nevertheless, empirical results indicate that this training paradigm significantly restricts the degree to which the machine learned model can fit non-physical spurious patterns in the data. We leave for future work developing a model that can predict the entire relaxation path starting with the only the highest energy state.

Using the 296 training relaxation paths, we carry out the model fitting algorithm described in Sec. II C. For the training set, we randomly select according to a uniform distribution, with replacement, 296 relaxation paths from the training set. Those paths selected more than once are repeated in the training set with the number of copies equaling the number of times the path was selected. Those paths that are not selected are placed in the validation set. The sparse linear model is trained using the greedy OLS algorithm with randomized feature bagging, with the number of features  $M$  ranging from  $M = 1$  to  $M = M_{\max} = 1000$ . The optimal number of features  $M = M^*$  is selected by minimizing the loss on the validation set. This procedure is repeated 100 times, resulting in 100 sparse linear models of the form (12), which are averaged together to yield the final model.

This final model is evaluated on the withheld test set. Figure 5 depicts the training, validation, and testing errors as a function of the number of model features  $M$ . It indicates that best models have, generally, between 64 and 256 features, with an average of 121 features per model, a small number given that there are  $\sim 70\,000$  training structures. Furthermore, the validation curve closely follows the test curve, indicating that our cross-validation procedure is nearly optimal for these test data. The average root mean squared



**FIG. 5.** Errors on the amorphous  $\text{Li}_x\text{Si}$  database as a function of number of features included in the model on a log-log scale. Error on the training set is shown in red, the validation set is shown in green, and the test set is shown in blue. The training error is a decreasing function of the number of features, whereas the validation and testing curves are not. The value  $M^*$  that minimizes the validation curve is the algorithm's best estimate for the optimal model that best balances under- and over-fitting of the training data. It has good agreement with the minimum of the test error curve.

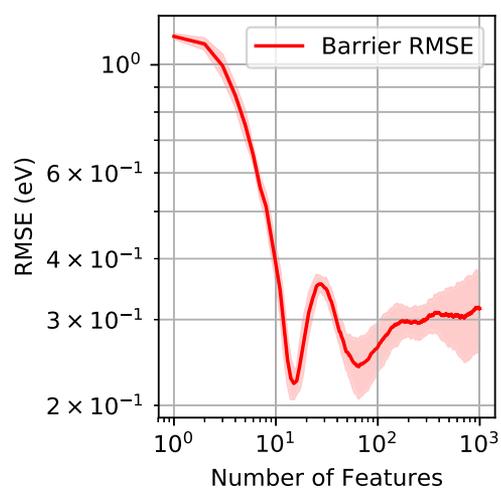
**TABLE I.** Numerical results for ML predictions on the test data from the amorphous dataset and the three extrapolation tasks from the model trained only on the amorphous data.

	RMSE (meV/atom)	MAE (meV/atom)
Relaxation paths	$7.44 \pm 0.49$	$5.52 \pm 0.34$
Diffusion	$12.3 \pm 0.50$	$11.7 \pm 0.51$
Large states	$9.54 \pm 0.25$	$6.81 \pm 0.23$
Bulk modulus	$12.8 \pm 1.36$	$8.92 \pm 0.68$

error (RMSE) and the average mean absolute error (MAE) over the five test folds, along with the standard deviation, are reported in the first row (relaxation paths) of Table I. Despite the small number of features, the RMSE is 7.44 meV/atom and the MAE is 5.52 meV/atom, which is comparable to the results reported by Onat *et al.*<sup>27</sup> and Artrith, Urban, and Ceder,<sup>25</sup> both of which used neural networks, and is small enough to be of use in materials science applications. However, the model developed here is significantly simpler than neural network models, being a linear model over multiscale, invariant features that utilize a universal set of filters. As such, the model is adept at generalization, as reported in Subsections III B and III C.

## B. Extrapolate: Diffusion in amorphous system

One important application of atomistic simulation is the study of atomic migration from site to site. The energetic barrier to migration determines diffusion constants and ionic conductivity.<sup>54</sup> The diffusion process may be simulated by directly tracking the mean square displacement using molecular dynamics or by calculating the migration barrier and using the Nernst-Einstein relationship. The first step in the explicit calculation is to find the minimum-energy path (MEP) for an atom to travel between two stable sites.



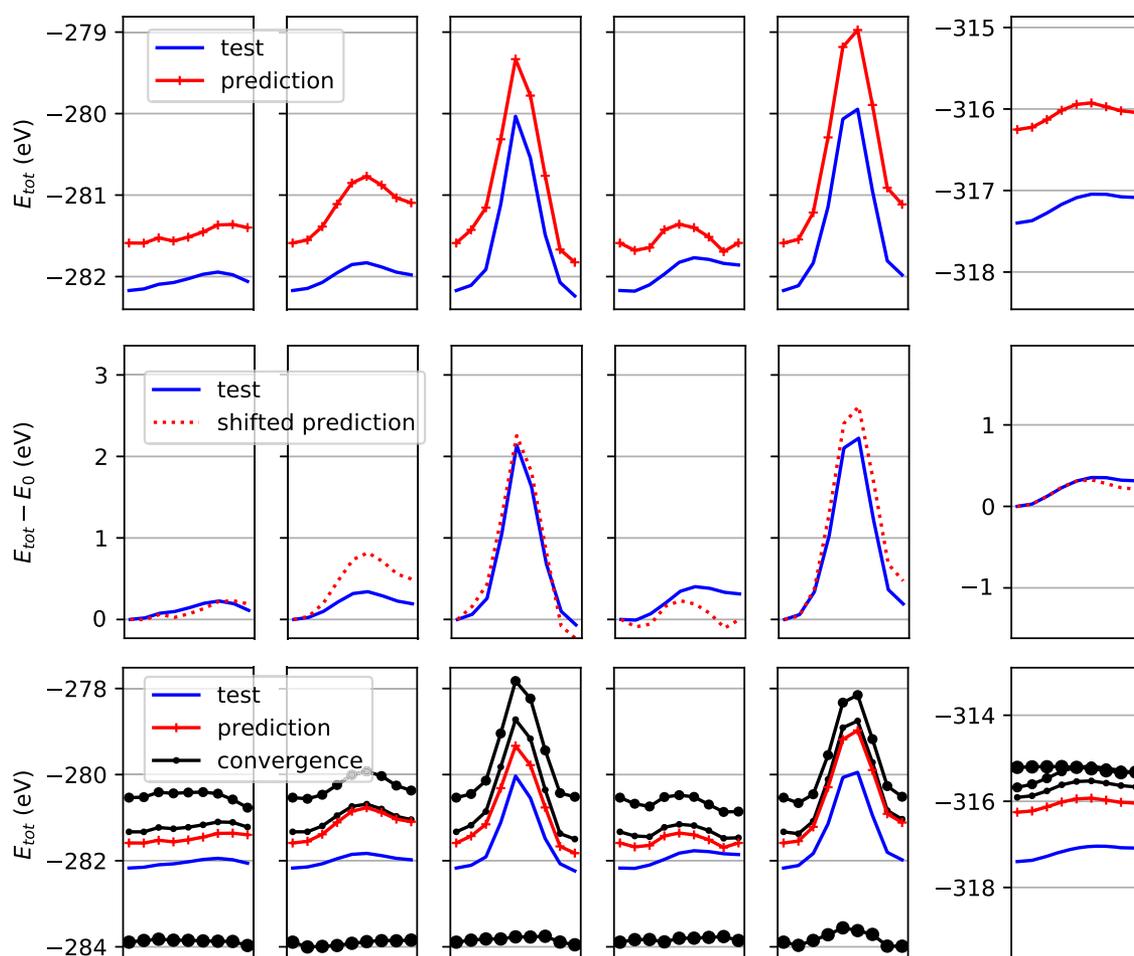
**FIG. 6.** Log-log plot of RMSE in diffusion barrier prediction averaged over the fivefolds.

This is typically done using optimization techniques such as the Nudged Elastic Band (NEB) method.<sup>51</sup> The barrier is defined as the energy difference between the stable position and the highest-energy position (saddle point) along the MEP.

There are a number of reasons why calculation of diffusion barriers may present a challenge for our ML model. Our present models do not predict forces, so they cannot be used with the NEB for prediction of the path itself. We therefore simply predict energies along the DFT-calculated MEP. A more fundamental challenge is the fact that the transition state structure, with one atom in a high-energy state and the rest in relatively low-energy states, is qualitatively different from the training items in the amorphous  $\text{Li}_\alpha\text{Si}$  dataset. Calculation of diffusion barrier is thus an extrapolation task. Furthermore, there is only one diffusing atom in the simulation box during calculation of the diffusion barrier. This means that energy per atom is no longer the most relevant measure of error. Instead, total energy

differences between similar structures along the MEP are the relevant quantity. Cancellation of systematic errors in DFT allows the calculation of energy differences along diffusion paths with much higher accuracy than would be suggested based on the accuracy of the method in total energy per atom.<sup>54</sup> It remains to be seen if similar cancellation of errors can improve the accuracy of diffusion barriers predicted by an ML model.

To test the extrapolation of our model to diffusion barriers, void spaces were identified in  $\text{Li}_{0.2}\text{Si}$  and  $\text{Li}_{0.5}\text{Si}$  by Voronoi analysis. Candidate endpoint structures were created by moving nearby lithium atoms into the voids and relaxing the resulting structure while keeping the target lithium atom fixed. Six endpoints were identified in which the void space was a local optimum for the lithium atom and in which the relaxation for the rest of the structure was minimal. These endpoints were then used together with the original  $\text{Li}_\alpha\text{Si}$  structures as the basis for NEB calculations. The structures



**FIG. 7.** Plots of the six diffusion barrier paths (blue) and (top row) model predictions in red, (middle row) model predictions and test data shifted by their respective starting-point energies  $E_0$ , and (bottom row) convergence of models with increasing number of features used for predictions of diffusion barrier curves. The large radii circles coincide with fewer features used starting from a model with a single feature. The models quickly converge in shape and progress toward the red curve, which is the aggregate model prediction. There is a curve for each choice of number of features  $M \in \{1, 21, 41\}$ .

**TABLE II.** Diffusion barriers (in eV) along various paths as predicted by our ML model and DFT. Paths 1–5 start from the same  $\text{Li}_{0.2}\text{Si}$  structure and path 6 is in  $\text{Li}_{0.5}\text{Si}$ .

Path	Barrier (ML model)	Barrier (DFT)
1	0.228	0.226
2	0.819	0.341
3	2.256	2.139
4	0.230	0.402
5	2.613	2.224
6	0.326	0.354

along the resulting NEB path were then passed to the ML model for comparison with the DFT results.

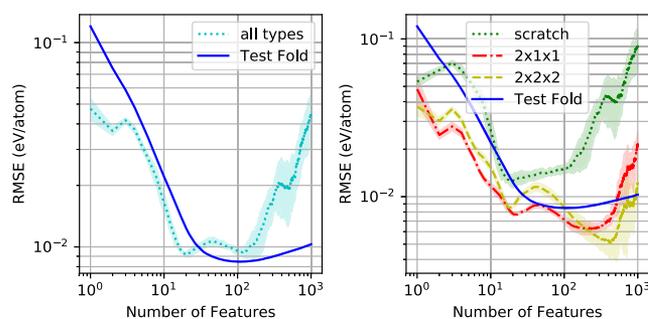
The learning curves are shown in Fig. 6. The RMSE for the diffusion path structures is less smooth than the RMSE for test folds consisting of the relaxation paths in the amorphous  $\text{Li}_a\text{Si}$  data. Table I, second row (diffusion), shows the RMSE and MAE of the per atom energy across all diffusion barrier structures. The RMSE for these structures is about 12.3 meV/atom, which is worse than on the relaxation path test but by less than a factor of two. Nevertheless, reduced accuracy is expected given the extrapolative nature of the task.

However, these errors are not the diffusion barrier errors, which is the quantity of interest. The energies along the diffusion paths are shown in Fig. 7. The first row plots the absolute energies for both the DFT calculation and the model prediction. The second row shifts the DFT and predicted energy curves to both start at zero, to more easily compare and read off the barriers, which are given in Table II. The third row of Fig. 7 plots the predicted energy curves as a function of the number of model features  $M$ , showing the learning rate of the model with respect to this task. The plots indicate that even with a small number of features, for example,  $M = 21$  or  $M = 41$ , the energy curve and resulting barrier is qualitatively correct, with additional features serving to refine the curves and better align the total energies.

Visual inspection of the energy along the diffusion paths shows that much of the error is systematic. The  $\text{Li}_{0.2}\text{Si}$  structures contain 60 atoms, so 12.3 meV/atom corresponds to 0.74 eV in total energy. If these errors were random, we would expect at least 0.74 eV error in prediction of the diffusion barrier. However, the curves show that the ML model can successfully distinguish between small-barrier paths and large-barrier paths, and the MAE in barrier prediction is 0.20 eV. While there is certainly room for improvement, we believe that these data show evidence that the ML model is able to partially capture the physics involved in the diffusion process.

### C. Extrapolate: Larger amorphous systems

It is desirable for an energy-predictor to generalize to structures in simulation cells with a different size than the training set so that it can be applied to simulation cells large enough to contain the geometries of experimental interest. As system size increases, the computation becomes challenging to carry out with DFT, but the wavelet scattering transform and linear regression scales efficiently with system size (for more details, see Appendix B), and we are thus much less inhibited by large systems.



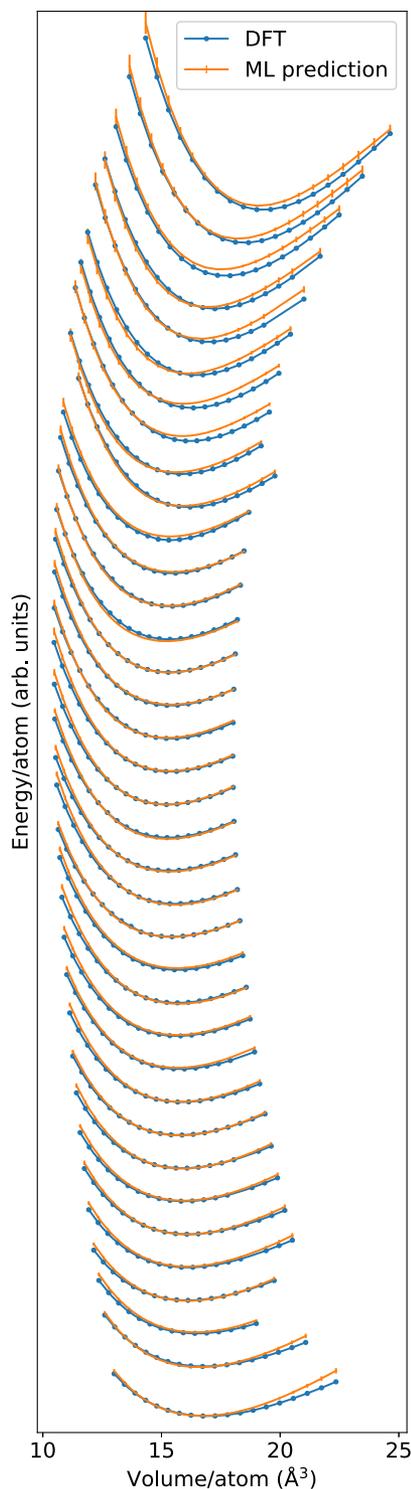
**FIG. 8.** A log-log plot of average of RMSEs of models on the interpolation test set from Sec. III A and on all types of large states (scratch,  $2 \times 1 \times 1$ ,  $2 \times 2 \times 2$ ) from Sec. III C. Here, y-axis =  $\log(\text{eV/atom})$  and x-axis =  $\log(\text{number of features in models})$ . The curves labeled  $2 \times 1 \times 1$ ,  $2 \times 2 \times 2$ , and scratch on right are the RMSE of energy error predictions of the five aggregate models separated by test folds. On the left panel, we see that the location of the minimum (i.e., the optimal number of features) for the interpolation test error is similar to the optimal number of features for the extrapolation error on larger states, although model over-fitting is significantly more costly for the larger states' predictions.

As discussed in Sec. II B, the data for this task were generated by two different methods: “from scratch” and “tiled.” The learning curves for each are shown in the right panel of Fig. 8. Since our model predicts global energies per atom, it gives the exact same result for a system that is simply periodically duplicated. This suggests that the predictions made when extrapolating to tiled systems that have been perturbed should maintain reasonable accuracy. This figure agrees with this conjecture since the corresponding error lines follow a similar trajectory to the line for the small system test data. This figure also shows that simpler models are favored for the independently relaxed AIMD-generated systems (the “from scratch” systems). These systems are less likely than the tiled systems to be similar to examples from the training set. The rapid increase in error on large systems for higher model complexity illustrates the sensitivity of the task to over-fitting. Nevertheless, as depicted in the left panel of Fig. 8, the optimal number of features for interpolation on amorphous  $\text{Li}_a\text{Si}$  data is approximately the same as the optimal number of features for energy predictions on the collection of states with larger unit cells. From Table I (third row, “large states”), we see that while the prediction errors are higher for the larger systems, it is not an unreasonable increase from errors on the smaller systems.

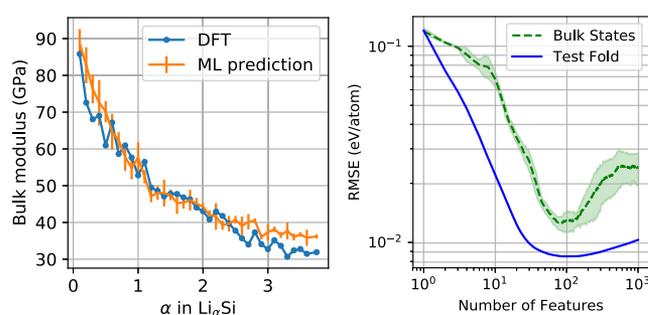
### D. Extrapolate: Bulk modulus

Elastic properties are another important output of atomistic simulations. These are typically calculated by applying small strains to the system in question and fitting elastic constants to the energy-vs-strain.<sup>54</sup> This too is an extrapolation task for our model because uniformly expanded or compressed structures do not appear in the training set. Testing data for this task were generated based on the lowest-energy structure at each concentration by applying hydrostatic strain and varying the side-length of the simulation box from  $-9\%$  to  $9\%$ .

The energy vs volume of the strained structures (Fig. 9) shows remarkable agreement between DFT and the ML model. The RMSE



**FIG. 9.** Energy per atom of hydrostatically strained  $\text{Li}_\alpha\text{Si}$  structures as a function of volume per atom. Energies are shifted vertically to avoid overlap:  $\alpha$  increases down the vertical axis. Error bars on ML prediction show the standard deviation of predictions of the fivefold cross-validated models for each structure.



**FIG. 10.** (Left) Comparison of DFT-calculated bulk modulus and ML-predicted bulk modulus. Modulus was calculated through a parabolic fit to points within  $\pm 4\%$  strain of the energy minimum. Error bars on the ML prediction show the standard deviation of fitted modulus across the fivefold cross-validated models. Averaging across the folds leads to a prediction with an MAE of 3.3 GPa compared to the DFT values. (Right) A log-log plot of average of RMSEs of models on the interpolation test set from Sec. III A and on bulk modulus data from Sec. III D. Here, y-axis =  $\log(\text{eV/atom})$  and x-axis =  $\log(\text{number of features in models})$ . Green curve is the average of the RMSEs for each fold with error bar given by the standard deviation over the fivefolds. As for the large states (see Fig. 8), we see that the location of the minimum (i.e., the optimal number of features) for the interpolation test error is similar to the extrapolation error for the bulk modulus states.

curves shown in Fig. 10 and the average errors in the last row of Table I (bulk modulus) are also quite low. The predicted bulk modulus of the structures is shown to decrease as a function of lithium content in Fig. 10. The ML method accurately captures lithiation-induced softening of the silicon. Energy-vs-strain curves along different deformation paths may also be used for the estimation of additional thermodynamic parameters, including Young's modulus, shear modulus free energies, and heat capacities through the Debye-Grüneisen model.<sup>55,56</sup>

#### IV. CONCLUSION

We have demonstrated a machine-learning model based on atomic orbital wavelet scattering that can achieve an accuracy of 5.52 meV/atom (mean absolute error) in energy on the prediction of amorphous  $\text{Li}_\alpha\text{Si}$  structures. We have tested the generalizability of this energy predictor on three extrapolation tasks: diffusion barriers, large systems, and bulk moduli. As expected based on the nature of regression-based ML, if care is not taken to avoid over-fitting, the model performs poorly on these extrapolation tasks. However, we have shown that a statistically based feature randomization procedure, using the universal wavelet scattering features, can significantly enhance performance on the extrapolation tasks without significant reduction in performance on the interpolative test set.

Although the present work is limited to amorphous  $\text{Li}_\alpha\text{Si}$ , it provides general lessons for those wishing to apply ML models to new problems in chemical physics. This is often a daunting task because ML is generally an interpolative technique. Before a model can be used, it must be trained on large amounts of data similar to the task at hand. If the problem is new or challenging to solve by conventional means, the generation of these data can be quite difficult. Extrapolation from the well-known systems may be possible, but

off-the-shelf ML models do not extrapolate well. However, extrapolation performance can be greatly improved by taking a different approach to training the ML model.

Simpler models generalize better. In our model, “simplicity” corresponds to the number of features (wavelet scattering coefficients) used and the fact that these features provide unsupervised descriptions of atomic states, but the concept is general. Validation sets are often used in machine learning to choose a model complex enough to describe the training data but simple enough to avoid over-fitting. By utilizing randomized feature selection and the aggregation of an ensemble of models (bootstrapping), we obtain a robust and accurate model when applied to the aforementioned extrapolation tasks. From this perspective, typical ML metrics such as testing and validation error are not the only criteria for a “good” model.

In order to apply these principles to harder extrapolation tasks and to incorporate *a priori* uncertainty quantification, it will be necessary to leverage statistical methods that allow one to predict which properties will be difficult for the model, suggesting possibilities for efficient training set expansion to further improve generalizability. Training set expansion could be automated using “active learning,” allowing a model to improve itself based on problems presented to it. The linear regression model over unsupervised nonlinear wavelet scattering features is well positioned for such future work, as it is relatively simple (compared to fully supervised neural networks) to incorporate new data on the fly.

In future work, we will extend the model to include force predictions. As described in (12), our energy predictions are given as linear combinations of features that are each dependent on the atomic positions. The features are differentiable and we can carry out this differentiation analytically, which opens up fast force computations and fits the weights of our model to force data. In this case, all the methods of Sec. II C would still be applicable. This has a computational advantage over features with learned filters, which would likely use automatic differentiation. Including forces in the weight learning process could affect the weights that are learned for the energy predictions as each system would have  $3N_x$  more points of training data. We expect that this will boost the generalization ability of the model. In this regard, training on forces will act as a regularizer for energy predictions.

## SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for voltages vs Li/Li<sup>+</sup> and radial distribution functions for Si–Si, Li–Si, and Li–Li pairs.

## AUTHORS' CONTRIBUTIONS

P.S. and M.W.S. contributed equally to this work.

## ACKNOWLEDGMENTS

This work was supported, in part, by the Defense Advanced Research Projects Agency (Young Faculty Award No. D16AP00117 to M.H., supporting P.S., X.B., J.L., and K.J.K.), the Alfred P. Sloan Foundation (Sloan Fellowship Grant No. FG-2016-6607 to M.H.), the National Science Foundation (Grant No. 1832808 to Y.Q., Grant No. 1620216 to M.H., and CAREER Award No. 1845856 to M.H.),

and the Michigan State University Foundation (Strategic Partnership Grant to Y.Q.). This work used computational resources provided by the Institute for Cyber-Enabled Research at Michigan State University. X.B. also acknowledges support from the Institute for Pure and Applied Mathematics at UCLA, where he was in residence during fall 2019.

## APPENDIX A: CHANNEL DEFINITIONS

In this appendix, we give precise definitions of the five input channels:

- Lithium channel:  $f(Z_k) = Z_k$  if  $Z_k = 3$  and  $f(Z_k) = 0$  otherwise.
- Silicon channel:  $f(Z_k) = Z_k$  if  $Z_k = 14$  and  $f(Z_k) = 0$  otherwise.
- Valence channel:  $f(Z_k) = \text{No. of valence electrons}$ .
- Ionic channel:  $f(Z_k) = \text{No. of core electrons}$ .
- Kinetic channel:  $f(Z_k) = \sqrt{Z_k}$ .

## APPENDIX B: FAST WAVELET SCATTERING COMPUTATIONS

In this appendix, we describe how to efficiently compute the wavelet scattering features described in Sec. II D.

In practice, all computations are carried out over a discrete sampling of the unit cell  $Q_x$ . Such a sampling is a three-dimensional grid  $G_x \subset Q_x$  with  $L_x$  grid points along each dimension. Due to the multiscale sizes of the wavelet filters  $\psi_{j,n,\ell}^m$ , a direct computation of the circular convolution (8) over the grid  $G_x$  will require  $O(L_x^6)$  floating point operations. This computational cost can be significantly reduced by carrying out these computations in frequency space.

Recalling (9), the Fourier transform of  $\rho_x^f * \psi_{j,n,\ell}^m$  is

$$\mathcal{F}[\rho_x^f * \psi_{j,n,\ell}^m](\omega) = \widehat{\psi}_{n,\ell}^m(2^j \omega) \sum_{k=1}^{N_x} f(Z_k) e^{-i\omega \cdot R_k}, \quad (\text{B1})$$

where  $\mathcal{F}[h](\omega) = \widehat{h}(\omega)$  is the Fourier transform of the function  $h \in \mathbf{L}^1(\mathbb{R}^3)$ . The Fourier transform of  $\psi_{n,\ell}^m$  can be computed analytically,

$$\widehat{\psi}_{n,\ell}^m(\omega) = (-i)^\ell \sqrt{\frac{4\pi}{2\ell+1}} |\omega|^{2(n-1)-\ell} e^{-\beta^2 |\omega|^2 / 2} Y_\ell^m(\omega/|\omega|).$$

Therefore, (B1) can be evaluated directly for any  $\omega \in \mathbb{R}^3$ . We do so in a box  $[-\omega_x, \omega_x]^3$ , where

$$\omega_x = \frac{\pi}{\Delta_x} \quad \text{and} \quad \Delta_x = \frac{s_x}{L_x}.$$

The maximum frequency  $\omega_x$  is chosen so that the essential support of  $\widehat{\psi}_{n,\ell}^m$  is contained within  $[-\omega_x, \omega_x]^3$ , which, in turn, determines the number of grid points  $L_x$  along each side length of the unit cell  $Q_x$ . This maximum frequency depends on the wavelet width  $\beta$  and (weakly) on the  $(n, \ell)$  parameters. Evaluations within the box  $[-\omega_x, \omega_x]^3$  are restricted to a grid  $\Omega_x \subset [-\omega_x, \omega_x]^3$  with grid spacing  $2\pi/s_x$ , which yields  $L_x$  frequency grid points. In

particular, we compute, via direct numerical evaluation, a tensor  $\Psi_{j,n,\ell}^m \in \mathbb{C}^{L_x} \times \mathbb{C}^{L_x} \times \mathbb{C}^{L_x}$  defined as

$$\Psi_{j,n,\ell}^m = \mathcal{F}[\rho_x^f * \psi_{j,n,\ell}^m] \Big|_{\Omega_x}. \quad (\text{B2})$$

Due to the discretization in (B2), taking the inverse fast Fourier transform (iFFT) of  $\Psi_{j,n,\ell}^m$  recovers the circular convolution  $\rho_x^f \otimes \psi_{j,n,\ell}^m$  evaluated on the spatial grid  $\mathcal{G}_x$ ,

$$\text{iFFT}(\Psi_{j,n,\ell}^m) = \rho_x^f \otimes \psi_{j,n,\ell}^m \Big|_{\mathcal{G}_x}. \quad (\text{B3})$$

The direct computation of  $\Psi_{j,n,\ell}^m$  requires  $CN_x L_x^3$  floating point operations, whereas the iFFT calculation requires  $CL_x^3 \log L_x$  floating point operations. Therefore, the total cost is reduced to  $O((N_x + \log L_x)L_x^3)$ .

First order wavelet scattering features are estimated by applying the pointwise nonlinear operator  $\sigma$  to (B3) and estimating the  $L^q(\mathcal{Q}_x)$  integrals with a Riemann sum approximation. Second order wavelet scattering features are computed by taking the fast Fourier transform (FFT) of  $\sigma(\rho_x^f \otimes \psi_{j,n,\ell}^m) \Big|_{\mathcal{G}_x}$  and computing the second circular wavelet convolution via frequency multiplication with a direct evaluation of  $\widehat{\psi}_{n_2,\ell_2}^{m_2}(2^b \omega)$  on the grid  $\omega \in \Omega_x$ , followed by another iFFT, application of  $\sigma$ , and Riemann sum. The cost of each second order feature, given that (B3) must already be computed for the first order features, is  $O(L_x^3 \log L_x)$ .

### APPENDIX C: ALTERNATIVE MODELS COMPARISON

In this appendix, we describe two models similar to the one used in the main body of this text and compare the results.

The model used in the main body (hereafter referred to as the full model) has numerical results on the test set summarized in Table I, and the training method is described in Sec. II C. The results of two alternative models on the various tasks of this work are listed in Tables III and IV. The test folds of relaxation paths and Li<sub>α</sub>Si states of the diffusion, large states, and bulk modulus states are identical in all three model comparisons.

The first alternative model (hereafter the 0–1 model) is trained identically to the full model with five test folds (the test folds are identical for both models) and 100 randomly selected sets of relaxation strings (with replacement) for training, but with only zero and first order wavelet scattering features available for selection in training. This results in a total of 321 features (with bias) to select from compared to 3741 in the full model. Note that at each step of the

**TABLE III.** Numerical results for ML predictions with only zero and first order features (compared to zero, first, and second in Table I) on the test data from the amorphous dataset and the three extrapolation tasks from the model trained only on the amorphous data.

	RMSE (meV/atom)	MAE (meV/atom)
Relaxation paths	8.04 ± 0.59	5.99 ± 0.39
Diffusion	11.8 ± 0.48	9.51 ± 0.95
Large states	14.0 ± 0.68	10.2 ± 0.42
Bulk modulus	39.5 ± 4.82	25.1 ± 2.92

**TABLE IV.** Numerical results for ML predictions with the non-randomized model. The models are trained without random feature selection at each step of the greedy OLS algorithm, i.e., at each step, all features are available for selection.

	RMSE (meV/atom)	MAE (meV/atom)
Relaxation paths	7.50 ± 0.39	5.64 ± 0.28
Diffusion	11.6 ± 1.01	11.0 ± 1.03
Large states	9.78 ± 1.98	6.60 ± 0.81
Bulk modulus	16.6 ± 4.91	11.5 ± 3.55

greedy OLS training, the best feature is chosen from  $\sqrt{321} \approx 17$  features that are randomly selected from the remaining unselected features compared to  $\sqrt{3741} \approx 61$  in the full model. The model size  $M^*$  averaged over all 500 permutations of the training data is 121 in the full model and 108 for the 0–1 model, with standard deviations of 64 and 38, respectively. The numerical results for the 0–1 model are listed in Table III. The performance is comparable on the relaxation paths. On the diffusion states, the 0–1 model has slightly better performance in RMSE and MAE, but the standard deviation in MAE across the fivefolds is nearly double the full model. Furthermore, inspection of the barriers computed by the 0–1 model reveals that they are, in fact, slightly worse than the full model. The performance of the 0–1 model is significantly worse than the full model on the large and bulk states, again with a large spread in errors. This indicates that we get a statistically significant benefit by including second order features in the models.

The second alternative model (hereafter the non-randomized model) has the same features available as the full model and the same five test folds as the prior two models. The training set is randomly partitioned into four equally sized sets (selection by relaxation strings) with a model trained for each selection of a set as validation and the remaining three for training (i.e., nested fivefold cross-validation, as in Ref. 57). This results in four trainings for each test fold for a total of 20 models trained compared to the 500 trainings (five test sets with 100 training/validation splits) of the full model. This non-randomized procedure ensures uniform representation of the strings in the training, validation, and testing folds. During training of the non-randomized model, the OLS algorithm seeks the next best feature at each step from all remaining features rather than randomly selecting a subset of features to choose from as in the prior two models. The average value of  $M^*$  is 153 for the non-randomized model with a standard deviation of 82 across the 20 trainings. The performance of this model is similar to the full model on relaxation paths but with a significantly larger spread of the errors between models on diffusion, large states, and bulk states. Thus, the chance of a catastrophic error is higher. Furthermore, the RMSE and MAE are significantly larger on the bulk states. This indicates that the model over-fit the training data and did not generalize as well to the extrapolation tasks. Randomized training in the full model appears to mitigate the possibility of over-fitting.

### DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

## REFERENCES

- <sup>1</sup>K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, "Machine learning for molecular and materials science," *Nature* **559**, 547–555 (2018).
- <sup>2</sup>F. Noé, A. Tkatchenko, K.-R. Müller, and C. Clementi, "Machine learning for molecular simulation," *Ann. Rev. Phys. Chem.* **71**, 361–390 (2020).
- <sup>3</sup>K. Muller, S. Mika, G. Ratsch, K. Tsuda, and B. Scholkopf, "An introduction to kernel-based learning algorithms," *IEEE Trans. Neural Networks* **12**, 181–201 (2001).
- <sup>4</sup>A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, "Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons," *Phys. Rev. Lett.* **104**(4), 136403 (2010).
- <sup>5</sup>M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," *Phys. Rev. Lett.* **108**, 058301 (2012).
- <sup>6</sup>J. E. Moussa, "Comment on 'fast and accurate modeling of molecular atomization energies with machine learning'," *Phys. Rev. Lett.* **109**, 059801 (2012).
- <sup>7</sup>A. P. Bartók, R. Kondor, and G. Csányi, "On representing chemical environments," *Phys. Rev. B* **87**(16), 184115 (2013).
- <sup>8</sup>G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, "Machine learning of molecular electronic properties in chemical compound space," *New J. Phys.* **15**, 095003 (2013).
- <sup>9</sup>S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, "Comparing molecules and solids across structural and alchemical space," *Phys. Chem. Chem. Phys.* **18**, 13754–13769 (2016).
- <sup>10</sup>A. Shapeev, "Moment tensor potentials: A class of systematically improvable interatomic potentials," *Multiscale Model. Simul.* **14**, 1153–1173 (2016).
- <sup>11</sup>S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, "Machine learning of accurate energy-conserving molecular force fields," *Sci. Adv.* **3**, e1603015 (2017).
- <sup>12</sup>F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, "Bypassing the Kohn-Sham equations with machine learning," *Nat. Commun.* **8**, 872 (2017).
- <sup>13</sup>S. Chmiela, H. E. Sauceda, K.-R. Müller, and A. Tkatchenko, "Towards exact molecular dynamics simulations with machine-learned force fields," *Nat. Commun.* **9**, 3887 (2018).
- <sup>14</sup>T. Bereau, R. A. DiStasio, A. Tkatchenko, and O. A. von Lilienfeld, "Non-covalent interactions across organic and biological subsets of chemical space: Physics-based potentials parametrized from machine learning," *J. Chem. Phys.* **148**, 241706 (2018).
- <sup>15</sup>I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- <sup>16</sup>J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.* **98**(4), 146401 (2007).
- <sup>17</sup>J. Behler, "Neural network potential-energy surfaces for atomistic simulations," *Chem. Modell.* **7**, 1–41 (2010).
- <sup>18</sup>K. T. Schütt, F. Arbabzadah, S. Chmiela, K.-R. Müller, and A. Tkatchenko, "Quantum-chemical insights from deep tensor neural networks," *Nat. Commun.* **8**, 13890 (2017); [arXiv:1609.08259](https://arxiv.org/abs/1609.08259).
- <sup>19</sup>K. T. Schütt, P.-J. Kindermans, H. E. Sauceda, S. Chmiela, A. Tkatchenko, and K.-R. Müller, "SchNet: A continuous-filter convolutional neural network for modeling quantum interactions," in *NIPS 2017* (Curran Associates, Inc., 2017), pp. 991–1001.
- <sup>20</sup>J. S. Smith, O. Isayev, and A. E. Roitberg, "ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost," *Chem. Sci.* **8**, 3192–3203 (2017).
- <sup>21</sup>J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *ICML 2017*, 2017.
- <sup>22</sup>T. S. Hy, S. Trivedi, H. Pan, B. M. Anderson, and R. Kondor, "Predicting molecular properties with covariant compositional networks," *J. Chem. Phys.* **148**, 241745 (2018).
- <sup>23</sup>L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, "Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics," *Phys. Rev. Lett.* **120**, 143001 (2018).
- <sup>24</sup>N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, "Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds," *NeurIPS Workshop on Machine Learning for Molecules and Materials*, Montreal, Canada, 2018. Workshop webpage: <http://www.quantum-machine.org/workshops/nips2018/>.
- <sup>25</sup>N. Artrith, A. Urban, and G. Ceder, "Constructing first-principles phase diagrams of amorphous Li<sub>x</sub>Si using machine-learning-assisted sampling with an evolutionary algorithm," *J. Chem. Phys.* **148**, 241711 (2018).
- <sup>26</sup>Z. Li, J. R. Kermode, and A. De Vita, "Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces," *Phys. Rev. Lett.* **114**, 096405 (2015).
- <sup>27</sup>B. Onat, E. D. Cubuk, B. D. Malone, and E. Kaxiras, "Implanted neural network potentials: Application to Li-Si alloys," *Phys. Rev. B* **97**, 094106 (2018).
- <sup>28</sup>T. Stecher, N. Bernstein, and G. Csányi, "Free energy surface reconstruction from umbrella samples using Gaussian process regression," *J. Chem. Theory Comput.* **10**, 4079–4097 (2014).
- <sup>29</sup>L. Mones, N. Bernstein, and G. Csányi, "Exploration, sampling, and reconstruction of free energy surfaces with Gaussian process regression," *J. Chem. Theory Comput.* **12**, 5100–5110 (2016).
- <sup>30</sup>E. Schneider, L. Dai, R. Q. Topper, C. Drechsel-Grau, and M. E. Tuckerman, "Stochastic neural network approach for learning high-dimensional free energy surfaces," *Phys. Rev. Lett.* **119**, 150601 (2017).
- <sup>31</sup>F. Noé, S. Olsson, J. Köhler, and H. Wu, "Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning," *Science* **365**, eaaw1147 (2019).
- <sup>32</sup>A. Voulozimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, "Deep learning for computer vision: A brief review," *Comput. Intell. Neurosci.* **2018**, 7068349.
- <sup>33</sup>A. Gatt and E. Kraemer, "Survey of the state of the art in natural language generation: Core tasks, applications and evaluation," *J. Artif. Intell. Res.* **61**, 65–170 (2018).
- <sup>34</sup>J. Hirschberg and C. D. Manning, "Advances in natural language processing," *Science* **349**, 261–266 (2015).
- <sup>35</sup>M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado, M. Hughes, and J. Dean, "Google's multilingual neural machine translation system: Enabling zero-shot translation," *Trans. Assoc. Comput. Linguistics* **5**, 339–351 (2017).
- <sup>36</sup>S. Mallat, "Group invariant scattering," *Commun. Pure Appl. Math.* **65**, 1331–1398 (2012).
- <sup>37</sup>X. Brumwell, P. Sinz, K. J. Kim, Y. Qi, and M. Hirn, "Steerable wavelet scattering for 3D atomic systems with application to Li-Si energy prediction," in *NeurIPS Workshop on Machine Learning for Molecules and Materials*, Montreal, Canada, 2018. Workshop webpage: <http://www.quantum-machine.org/workshops/nips2018/>.
- <sup>38</sup>G. Imbalzano, A. Anelli, D. Giofré, S. Klees, J. Behler, and M. Ceriotti, "Automatic selection of atomic fingerprints and reference configurations for machine-learning potentials," *J. Chem. Phys.* **148**, 241730 (2018).
- <sup>39</sup>R. Drautz, "Atomic cluster expansion for accurate and transferable interatomic potentials," *Phys. Rev. B* **99**, 014104 (2019).
- <sup>40</sup>M. Hirn, S. Mallat, and N. Poilvert, "Wavelet scattering regression of quantum chemical energies," *Multiscale Model. Simul.* **15**, 827–863 (2017).
- <sup>41</sup>M. Eickenberg, G. Exarchakis, M. Hirn, and S. Mallat, "Solid harmonic wavelet scattering: Predicting quantum molecular energy from invariant descriptors of 3D electronic densities," in *NIPS 2017* (Curran Associates, Inc., 2017), pp. 6540–6549.
- <sup>42</sup>M. Eickenberg, G. Exarchakis, M. Hirn, S. Mallat, and L. Thiry, "Solid harmonic wavelet scattering for predictions of molecule properties," *J. Chem. Phys.* **148**, 241732 (2018).
- <sup>43</sup>T. P. Sentfle, S. Hong, M. M. Islam, S. B. Klyasa, Y. Zheng, Y. K. Shin, C. Junkermeier, R. Engel-Herbert, M. J. Janik, H. M. Aktulga, T. Verstraelen, A. Grama, and A. C. T. van Duin, "The ReaxFF reactive force-field: Development, applications and future directions," *npj Comput. Mater.* **2**, 15011 (2016).
- <sup>44</sup>Y. Yuan, G. Houchins, P.-W. Guan, and V. Viswanathan, "Uncertainty quantification of first principles computational phase diagram predictions of Li-Si system via Bayesian sampling," [arXiv:2003.13393](https://arxiv.org/abs/2003.13393) [cond-mat.mtrl-sci] (2020).
- <sup>45</sup>M. W. Swift and Y. Qi, "First-principles prediction of potentials and space-charge layers in all-solid-state batteries," *Phys. Rev. Lett.* **122**, 167701 (2019).

- <sup>46</sup>S.-M. Liang, F. Taubert, A. Kozlov, J. Seidel, F. Mertens, and R. Schmid-Fetzer, "Thermodynamics of Li-Si and Li-Si-H phase diagrams applied to hydrogen absorption and Li-ion batteries," *Intermetallics* **81**, 32–46 (2017).
- <sup>47</sup>T. D. Hatchard and J. R. Dahn, "Study of the electrochemical performance of sputtered Si<sub>1-x</sub>Sn<sub>x</sub> films," *J. Electrochem. Soc.* **151**, A1628 (2004).
- <sup>48</sup>G. Kresse and J. Furthmüller, "Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set," *Phys. Rev. B* **54**, 11169–11186 (1996).
- <sup>49</sup>P. E. Blöchl, "Projector augmented-wave method," *Phys. Rev. B* **50**, 17953–17979 (1994).
- <sup>50</sup>J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximation made simple," *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- <sup>51</sup>H. Jónsson, G. Mills, and K. W. Jacobsen, "Nudged elastic band method for finding minimum energy paths of transitions," in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, 1998), p. 385.
- <sup>52</sup>F. Birch, "Finite elastic strain of cubic crystals," *Phys. Rev.* **71**, 809–824 (1947).
- <sup>53</sup>P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, "A high-bias, low-variance introduction to machine learning for physicists," *Phys. Rep.* **810**, 1–124 (2019).
- <sup>54</sup>D. Sholl and J. Steckel, *Density Functional Theory: A Practical Introduction* (Wiley, 2009).
- <sup>55</sup>X.-G. Lu, M. Selleby, and B. Sundman, "Calculations of thermophysical properties of cubic carbides and nitrides using the Debye–Grüneisen model," *Acta Mater.* **55**, 1215–1226 (2007).
- <sup>56</sup>P.-W. Guan, G. Houchins, and V. Viswanathan, "Uncertainty quantification of DFT-predicted finite temperature thermodynamic properties within the Debye model," *J. Chem. Phys.* **151**, 244702 (2019).
- <sup>57</sup>K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, "Assessment and validation of machine learning methods for predicting molecular atomization energies," *J. Chem. Theory Comput.* **9**, 3404–3419 (2013).