


Accelerating finite-temperature Kohn-Sham density functional theory with deep neural networksJ. A. Ellis,¹ L. Fiedler ,^{2,3} G. A. Popoola,¹ N. A. Modine,¹ J. A. Stephens,¹ A. P. Thompson,¹ A. Cangi,^{2,3,*} and S. Rajamanickam^{1,†}¹*Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*²*Center for Advanced Systems Understanding (CASUS), D-02826 Görlitz, Germany*³*Helmholtz-Zentrum Dresden-Rossendorf, D-01328 Dresden, Germany*

(Received 17 October 2020; revised 27 May 2021; accepted 21 June 2021; published 8 July 2021)

We present a numerical modeling workflow based on machine learning which reproduces the total energies produced by Kohn-Sham density functional theory (DFT) at finite electronic temperature to within chemical accuracy at negligible computational cost. Based on deep neural networks, our workflow yields the local density of states (LDOS) for a given atomic configuration. From the LDOS, spatially resolved, energy-resolved, and integrated quantities can be calculated, including the DFT total free energy, which serves as the Born-Oppenheimer potential energy surface for the atoms. We demonstrate the efficacy of this approach for both solid and liquid metals and compare results between independent and unified machine-learning models for solid and liquid aluminum. Our machine-learning density functional theory framework opens up the path towards multiscale materials modeling for matter under ambient and extreme conditions at a computational scale and cost that is unattainable with current algorithms.

DOI: [10.1103/PhysRevB.104.035120](https://doi.org/10.1103/PhysRevB.104.035120)**I. INTRODUCTION**

Multiscale materials modeling [1] provides fundamental insights into microscopic mechanisms that determine materials properties. A multiscale modeling framework operating both near first-principles accuracy and across length and time scales would enable key progress in a plethora of applications. It would greatly advance materials science research—in gaining understanding of the dynamical processes inherent to advanced manufacturing [2,3], in the search for superhard [4] and energetic materials [5,6], and in the identification of radiation-hardened semiconductors [7,8]. It would also pave the way for generating more accurate models of materials at extreme conditions including shock-compressed materials [9,10] and astrophysical objects, such as the structure, dynamics, and formation processes of the Earth's core [11], and solar [12–18] and extrasolar planets [19,20]. This requires faithfully passing information from quantum mechanical calculations of electronic properties on the atomic scale (nanometers and femtoseconds) up to effective continuum material simulations operating at much larger length and time scales. Atomistic simulations based on molecular dynamics (MD) techniques [21] and their efficient implementation [22] for large-scale simulations on high-performance computing platforms are the

key link. For atomic-scale systems, MD results can be directly validated against quantum calculations. At the same time, MD can be scaled up to the much larger scales on which the continuum behaviors start to emerge (either micrometers or microseconds, or even both [23]).

MD simulations require accurate interatomic potentials (IAPs) [24]. Generating them based on machine-learning (ML) techniques has only recently become a rapidly evolving research area. It has already led to several ML-generated IAPs such as AGNI [25], ANI [26], DPMD [27], GAP [28], GARFfield [29], HIP-NN [30], SchNet [31], and SNAP [32]. While these IAPs differ in their flavor of ML models, for instance, nonlinear optimization, kernel methods, structured neural networks, and convolutional neural networks, they all rely on accurate training data from first-principles methods.

First-principles training data sets are commonly generated with Kohn-Sham density functional theory (DFT) [33] and its generalization to finite electronic temperature [34]. Due to its balance between accuracy and computational cost, it is the method of choice [35] for computing properties of molecules and materials with close to chemical accuracy [36,37].

Despite its success [38], DFT is limited to the nanoscale due to its computational cost which scales as the cube of the system size. At the same time, the amount of data needed to construct an IAP increases exponentially with the number of chemical elements, thermodynamic states, phases, and interfaces [39]. Likewise, evaluation of thermodynamic properties using DFT is computationally expensive, typically 10^5 to 10^6 core-hours for each point on a grid of temperatures and densities [40]. Furthermore, fully converged DFT calculations at low densities, very high temperatures, or near phase transitions are considerably more difficult.

*a.cangi@hzdr.de

†srajama@sandia.gov

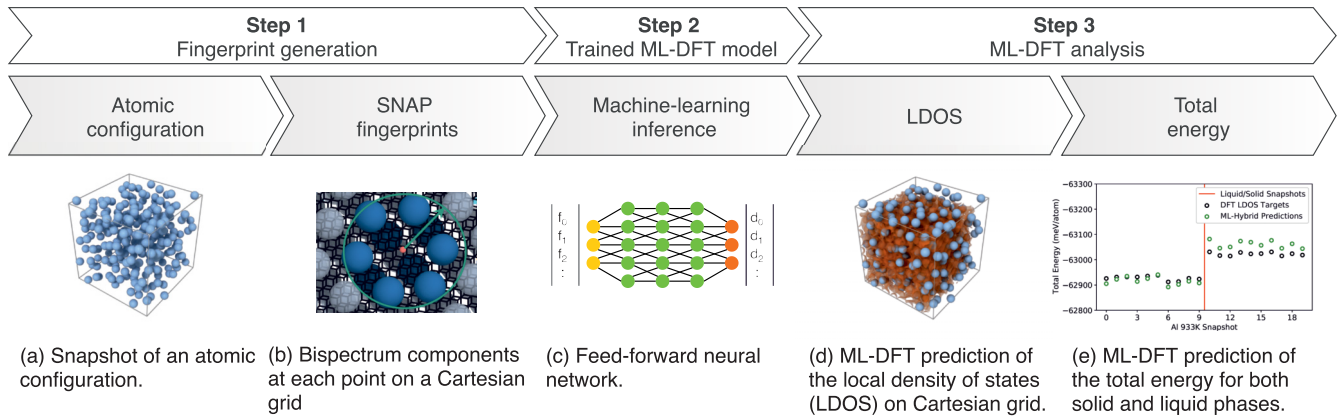


FIG. 1. Illustration of the ML-DFT workflow listed in Table I. (a) The raw input data are generated from finite-temperature DFT calculations for snapshots of atomic configurations, i.e., positions and chemical identities of atoms. (b) SNAP fingerprints in terms of bispectrum components are generated for each of several million Cartesian grid points in the snapshot volume. (c) The trained ML-DFT feed-forward neural network takes the input vector \mathbf{f} composed of SNAP fingerprints for one Cartesian grid point and outputs the predicted LDOS vector \mathbf{d} at the grid point. (d) The primary output of the ML inference is the LDOS on a Cartesian grid which is used to compute (e) the total energy based on Eq. (9). The embedded Fig. 6 displays our central result—the capability of a single ML-DFT model to predict the total energy in both solid and liquid phases of aluminum with accuracy comparable to standard DFT.

Pioneering efforts in applying ML to electronic structure calculations focus on mining benchmark databases generated from experiments, quantum chemistry methods, and DFT in order to train ML models based on kernel ridge regression and feed-forward neural networks. These efforts enabled ML predictions of molecular properties [41–43] and crystal structures [44]. Only limited efforts have gone into actually speeding up DFT calculations by directly approximating the solutions of the Kohn-Sham differential equations through ML models. Prior work has focused on approximating the kinetic energy functional using ML regression [45–47] and deep-learning models [48]. Most recently, ongoing efforts have demonstrated the usefulness of ML kernel techniques on the electronic density [49] and feed-forward neural networks on the density of states (DOS) [44,50,51] and the local density of states (LDOS) [52] for creating computationally efficient DFT surrogates.

In this work, we develop ML-DFT, a surrogate for DFT calculations at finite electronic temperature using feed-forward neural networks. We derive a representation of the Born-Oppenheimer potential energy surface at finite electronic temperature (Sec. II A) that lends itself to an implementation within an ML workflow (Sec. II B). Using the LDOS as the central quantity, ML-DFT enables us to compute the DFT total free energy in terms of a single ML model (Sec. II B), applicable to coupled electrons and ions not only under ambient conditions but also in liquid metals and plasmas. For the sake of simplicity, we will refer to the DFT total free energy as defined in Eq. (6) as total energy. Utilizing a generalization of the SNAP bispectrum descriptors on a Cartesian grid (Sec. II E), the feed-forward neural network of ML-DFT (Sec. II F) is trained on DFT data (Sec. II C). Integrated quantities such as the band energy and the total energy of an atomic configuration at a given electronic temperature are in turn computed based on the LDOS predicted by ML-DFT. Using the LDOS in conjunction with the developed expression for total energy avoids the need to introduce another ML model

to map densities to kinetic energies [49] and goes significantly beyond prior ML efforts related to the LDOS [52]. ML-DFT yields DFT-level accuracy, but comes at a negligible computational cost by avoiding the $O(N^3)$ eigensolver step. We demonstrate this for aluminum (Sec. III). Trained on atomic configurations of aluminum at the melting point, a single ML-DFT model yields accurate results (LDOS, electronic density, band energy, and total energy) for both crystalline and liquid aluminum. Notably, band energies and total energies predicted by ML-DFT agree with the results of conventional DFT calculations to well within *chemical accuracy*, which is traditionally defined as 1 kcal/mol = 43.4 meV/atom. Hence, ML-DFT meets all requirements to become the backbone of the computational infrastructure needed for high-performance multiscale materials modeling of matter under ambient and extreme conditions (Sec. IV).

The ML-DFT workflow is illustrated in Fig. 1. It starts with input data generated from finite-temperature DFT calculations for a given atomic configuration [Fig. 1(a)]. SNAP fingerprints are generated for each of several million Cartesian grid points in the snapshot volume [Fig. 1(b)]. Given a SNAP fingerprint as input, the trained ML-DFT feed-forward neural network makes its prediction for the LDOS vector [Fig. 1(c)]. The primary output of the ML inference is the LDOS prediction at each grid point [Fig. 1(d)]. Based on the predicted LDOS, quantities such as the electronic density based on Eq. (10), the DOS based on Eq. (13), and the total energy based on Eq. (9) are computed [Fig. 1(e)].

II. METHODS

A. Born-Oppenheimer density functional theory at finite electronic temperature

A suitable theoretical framework for computing thermodynamic materials properties from first principles is within the scope of nonrelativistic quantum mechanics in the Born-Oppenheimer approximation [53]. Additionally we work

within a formalism that takes the electronic temperature into account. This becomes relevant when the scale of the electronic temperature becomes comparable to the first excitation energy of the electronic system which is particularly relevant for liquid metals and plasmas. In the given framework we, hence, assume that the electrons find their thermal equilibrium on a time scale that is small compared to that of ionic motion. The formal development and implementation of such methodologies that couple electron and ion dynamics is an area of active research [54–60].

More formally, consider a system of N_e electrons and N_i ions with collective coordinates $\underline{r} = \{\mathbf{r}_1, \dots, \mathbf{r}_{N_e}\}$ and $\underline{\mathbf{R}} = \{\mathbf{R}_1, \dots, \mathbf{R}_{N_i}\}$, where $\mathbf{r}_j \in \mathbb{R}^3$ refers to the position of the j th electron, while $\mathbf{R}_\alpha \in \mathbb{R}^3$ denotes the position of the α th ion of mass M_α and charge Z_α . The physics in this framework is governed by the Born-Oppenheimer Hamiltonian

$$\hat{H}^{BO}(\underline{r}; \underline{\mathbf{R}}) = \hat{T}^e(\underline{r}) + \hat{V}^{ee}(\underline{r}) + \hat{V}^{ei}(\underline{r}; \underline{\mathbf{R}}) + \hat{V}^{ii}(\underline{\mathbf{R}}), \quad (1)$$

where $\hat{T}^e(\underline{r}) = \sum_j^{N_e} -\nabla_j^2/2$ denotes the kinetic energy of the electrons, $\hat{V}^{ee}(\underline{r}) = \sum_j^{N_e} \sum_{k \neq j}^{N_e} |\mathbf{r}_j - \mathbf{r}_k|^{-1}/2$ the electron-electron interaction, $\hat{V}^{ei}(\underline{r}; \underline{\mathbf{R}}) = -\sum_j^{N_e} \sum_\alpha^{N_i} Z_\alpha |\mathbf{r}_j - \mathbf{R}_\alpha|^{-1}$ the electron-ion interaction, and $\hat{V}^{ii}(\underline{\mathbf{R}}) = \sum_\alpha^{N_i} \sum_{\beta \neq \alpha}^{N_i} Z_\alpha Z_\beta |\mathbf{R}_\alpha - \mathbf{R}_\beta|^{-1}/2$ the ion-ion interaction. Note that we work within atomic units throughout, where $\hbar = m_e = e^2 = 1$, such that energies are expressed in hartrees and length in Bohr radii.

Using the standard framework of quantum statistical mechanics [61], all thermodynamic equilibrium properties are formally given in terms of the grand potential

$$\Omega[\hat{\Gamma}] = \text{Tr}(\hat{\Gamma} \hat{\Omega}), \quad (2)$$

which is defined as a statistical average over the grand canonical operator $\hat{\Omega} = \hat{H}^{BO} - \mu \hat{N} - \hat{S}/\beta$, where \hat{N} denotes the particle number operator, \hat{S} the entropy operator, μ the chemical potential, and $1/\beta = k_B T$, where k_B is the Boltzmann constant and T is the electronic temperature. Statistical averages as in Eq. (2) are computed via the statistical density operator $\hat{\Gamma} = \sum_{N_e, j} w_{N_e, j} |\Psi_{N_e, j}\rangle \langle \Psi_{N_e, j}|$ which is a weighted sum of projection operators on the underlying Hilbert space spanned by the N_e -particle eigenstates $\Psi_{N_e, j}$ of \hat{H}^{BO} with energies $E_{N_e, j}$ and statistical weights $w_{N_e, j} = \exp[-\beta(E_{N_e, j} - \mu N_e)] / \sum_{N_e, j} \exp[-\beta(E_{N_e, j} - \mu N_e)]$. The thermal equilibrium of a grand canonical ensemble is then defined as that statistical density operator which minimizes the grand potential. Further details on the grand canonical ensemble formulation of a thermal electronic system can be found in Refs. [62,63].

In practice, the electron-electron interaction complicates finding solutions to $\Psi_{N_e, j}$ and, hence, evaluating the grand potential as defined in Eq. (2). Instead, a solution is found in a computationally feasible manner within DFT [33] at finite electronic temperature [34,62,64]. Here, a formally exact representation of the grand potential is given by

$$\begin{aligned} \Omega[n] = & T_S[n] - S_S[n]/\beta + U[n] + E_{XC}[n] \\ & + \int d\mathbf{r} n(\mathbf{r}; \underline{\mathbf{R}}) [v^{ei}(\mathbf{r}; \underline{\mathbf{R}}) - \mu] + V^{ii}(\underline{\mathbf{R}}), \end{aligned} \quad (3)$$

where T_S denotes the Kohn-Sham kinetic energy, S_S the Kohn-Sham entropy, $U[n]$ the classical electrostatic interaction energy, E_{XC} the exchange-correlation free energy which is the sum of the temperature-dependent exchange-correlation energy and the interacting entropy, $v^{ei}(\mathbf{r}; \underline{\mathbf{R}}) = -\sum_\alpha Z_\alpha / |\mathbf{r} - \mathbf{R}_\alpha|$ the electron-ion interaction, and V^{ii} the ion-ion interaction. The grand potential in Eq. (3) is evaluated in terms of the electronic density defined by

$$n(\mathbf{r}; \underline{\mathbf{R}}) = \sum_j f^\beta(\epsilon_j) |\phi_j(\mathbf{r}; \underline{\mathbf{R}})|^2, \quad (4)$$

where the sum runs over the Kohn-Sham orbitals ϕ_j and eigenvalues ϵ_j that are obtained from the solving the Kohn-Sham equations

$$\left[-\frac{1}{2} \nabla^2 + v_S(\mathbf{r}; \underline{\mathbf{R}}) \right] \phi_j(\mathbf{r}; \underline{\mathbf{R}}) = \epsilon_j \phi_j(\mathbf{r}; \underline{\mathbf{R}}), \quad (5)$$

where $f^\beta(\epsilon_j) = (1 + \exp[\beta(\epsilon_j - \mu)])^{-1}$ denotes the Fermi-Dirac distribution. The Kohn-Sham framework is constructed such that a noninteracting system of fermions yields the same electronic density as the interacting many-body system of electrons defined by \hat{H}^{BO} . This is achieved by the Kohn-Sham potential $v_S(\mathbf{r}; \underline{\mathbf{R}}) = \delta U[n]/\delta n(\mathbf{r}; \underline{\mathbf{R}}) + \delta E_{XC}[n]/\delta n(\mathbf{r}; \underline{\mathbf{R}}) + v^{ei}(\mathbf{r}; \underline{\mathbf{R}})$ which includes all electron-electron interactions on a mean-field level. While formally exact, approximations to E_{XC} are applied in practice either at the electronic ground state [38] or at finite temperature [65–67].

The key quantity connecting the electronic and ionic degrees of freedom is the total energy

$$A^{BO}[n](\underline{\mathbf{R}}) = \Omega_0[n] + \mu N_e. \quad (6)$$

This expression yields the Born-Oppenheimer potential energy surface at finite electronic temperature, when it is evaluated as a function of $\underline{\mathbf{R}}$. It provides the forces $\mathbf{F}_\alpha = -\partial A^{BO}(\underline{\mathbf{R}})/\partial \mathbf{R}_\alpha$ on the ions which are obtained from solving for the instantaneous, grand potential in thermal equilibrium $\Omega_0[n]$ for a given atomic configuration $\underline{\mathbf{R}}$ via Eq. (5). In doing so, Eq. (6) enables us to time-propagate the dynamics of the atomic configuration $\underline{\mathbf{R}}$, in its simplest realization, by solving a Newtonian equation of motion,

$$M_\alpha \frac{d^2 \mathbf{R}_\alpha}{dt^2} = \mathbf{F}_\alpha. \quad (7)$$

This can be further extended, for example, to take into account the energy transfer between the coupled system of electrons and ions by introducing thermostats [68,69].

Despite its powerful utility for predicting thermodynamic and material properties on the atomic scale [36,37] and providing benchmark data for the parametrization of IAPs for use in MD simulations [29], applying DFT becomes computationally infeasible in applications towards the mesoscopic scale due to the computational scaling of Eq. (5). The computational cost for data generation explodes exponentially with the number of chemical elements, thermodynamic states, phases, and interfaces [39] needed for the construction of IAPs.

B. Machine-learning model of the total energy

We now reformulate DFT in terms of the LDOS which is granted by virtue of the first Hohenberg-Kohn theorem [70]. We then replace the DFT evaluation of the LDOS with a ML model for the LDOS in order to obtain all of the quantities required to evaluate the total energy while avoiding the computationally expensive solution of Eq. (5). The LDOS is defined as a sum over the Kohn-Sham orbitals,

$$D(\epsilon, \mathbf{r}; \mathbf{R}) = \sum_j |\phi_j(\mathbf{r}; \mathbf{R})|^2 \delta(\epsilon - \epsilon_j), \quad (8)$$

where ϵ denotes the energy as a continuous variable, and $\delta(\epsilon)$ is the Dirac delta function. This enables us to rewrite Eq. (6) as

$$A^{BO}[D](\mathbf{R}) = E_b[D] - S_S[D]/\beta - U[D] + E_{XC}[D] - V_{XC}[D] + V^{ii}(\mathbf{R}), \quad (9)$$

where $V_{XC}[D] = \int d\mathbf{r} n[D](\mathbf{r}; \mathbf{R}) \delta E_{XC}[D]/\delta n[D](\mathbf{r}; \mathbf{R})$ denotes the potential energy component of the exchange-correlation (XC) energy [71].

The central advantage of this reformulation is that, by virtue of Eq. (9), the total energy is expressed solely in terms of the LDOS. We emphasize this by explicitly denoting the functional dependence of each term on the LDOS.

Now let us turn to the individual terms in this expression. The ion-ion interaction $V^{ii}(\mathbf{R})$ is given trivially by the atomic configuration. The Hartree energy $U[D] = U[n[D]]$, the XC energy $E_{XC}[D] = E_{XC}[n[D]]$, and the potential XC energy $V_{XC}[D] = V_{XC}[n[D]]$ are all determined implicitly by the LDOS via their explicit dependence on the electronic density:

$$n[D](\mathbf{r}; \mathbf{R}) = \int d\epsilon f^\beta(\epsilon) D(\epsilon, \mathbf{r}; \mathbf{R}). \quad (10)$$

Likewise, the band energy

$$E_b[D] = \int d\epsilon f^\beta(\epsilon) \epsilon \bar{D}[D](\epsilon; \mathbf{R}) \quad (11)$$

and the Kohn-Sham entropy

$$S_S[D] = - \int d\epsilon \{ f^\beta(\epsilon) \log[f^\beta(\epsilon)] + [1 - f^\beta(\epsilon)] \log[1 - f^\beta(\epsilon)] \} \bar{D}[D](\epsilon; \mathbf{R}) \quad (12)$$

are determined implicitly by the LDOS via their explicit dependence on the DOS:

$$\bar{D}[D](\epsilon; \mathbf{R}) = \int d\mathbf{r} D(\epsilon, \mathbf{r}; \mathbf{R}). \quad (13)$$

Equation (9) enables us to evaluate the total energy in terms of a single ML model that is trained on the LDOS and can predict the LDOS for snapshots unseen in training. This is advantageous, as it has been shown that learning energies from electronic properties, such as the density, is significantly more accurate [49] than learning them directly from atomic descriptors [45, 72, 73]. Furthermore, developing an ML model based on the LDOS instead of the density avoids complications of prior models. For example, accuracy was lost due to an additional ML model needed to map densities to kinetic energies or noise was introduced by taking gradients of the ML

kinetic energy models [46, 47, 74, 75]. The two formulations in Eqs. (6) and (9) are mathematically equivalent. However, Eq. (9) is more convenient in our ML-DFT workflow where we use the LDOS as the central quantity.

C. Data generation

Our initial efforts have focused on developing an ML model for aluminum at ambient density (2.699 g/cc) and temperatures up to the melting point (933 K). The training data for this model were generated by calculating the LDOS for a set of atomic configurations using the QUANTUM ESPRESSO electronic structure code [76–78]. The structures were generated from snapshots of equilibrated DFT-MD trajectories for 256 atom supercells of aluminum. We calculated LDOS training data for ten snapshots at room temperature, ten snapshots of the crystalline phase at the melting point, and ten snapshots of the liquid phase at the melting point.

Our DFT calculations used a scalar-relativistic, optimized norm-conserving Vanderbilt pseudopotential (Al.sr-pbesol.upf) [79] and the PBEsol exchange-correlation functional [80]. We used a 100-Ry plane-wave cutoff to represent the Kohn-Sham orbitals and a 400-Ry cutoff to represent densities and potentials. These cutoffs resulted in a $200 \times 200 \times 200$ real-space grid for the densities and potentials, and we used this same grid to represent the LDOS. The electronic occupations were generated using Fermi-Dirac smearing with the electronic temperature corresponding to the ionic temperature. Monkhorst-Pack [81] k -point sampling grids, shifted to include the Γ point when necessary, were used to sample the band structure over the Brillouin zone. The DFT-MD calculations that generated the snapshots used $2 \times 2 \times 2$ k -point sampling for the crystalline phase and $1 \times 1 \times 1$ k -point sampling for the liquid phase.

In contrast, the generation of the LDOS used $8 \times 8 \times 8$ k -point sampling. The reason for this unusually high k -point sampling (for a 256-atom supercell) was to help overcome a challenge that arises when using the LDOS as the fundamental parameter in a ML surrogate for DFT. For a periodic supercell, the LDOS is given as

$$D(\epsilon, \mathbf{r}; \mathbf{R}) = \Omega_{BZ}^{-1} \int_{BZ} d\mathbf{k} \sum_j^{N_s} |\phi_{jk}(\mathbf{r}; \mathbf{R})|^2 \delta(\epsilon - \epsilon_{jk}), \quad (14)$$

where the index j now labels the Kohn-Sham orbitals at a particular \mathbf{k} value, and the vector index \mathbf{k} categorizes the Kohn-Sham orbitals by their transformation properties under the group of translation symmetries of the periodic supercell. In particular,

$$\phi_{jk}(\mathbf{r}; \mathbf{R}) = u_{jk}(\mathbf{r}; \mathbf{R}) \exp(i\mathbf{k} \cdot \mathbf{r}), \quad (15)$$

where $u_{jk}(\mathbf{r}; \mathbf{R})$ is some function with the same periodicity as the periodically repeated atomic positions \mathbf{R} .

In Eq. (14), the integral with respect to \mathbf{k} is taken over the first Brillouin zone (BZ), which has volume Ω_{BZ} . If this integral could be evaluated exactly, the resulting LDOS would be a continuous function of ϵ with Van Hove singularities [82] of the form $\sqrt{\epsilon - \epsilon'}$ or $\sqrt{\epsilon' - \epsilon}$ at energies ϵ' that correspond to critical points of the band structure ϵ_{jk} . However, when this integral is represented as a summation over a finite number of

k points $\mathbf{k}_1, \dots, \mathbf{k}_{N_k}$, as in the Monkhorst-Pack [81] approach used in our calculations, it is necessary to replace the Dirac delta function $\delta(\epsilon - \epsilon_{j\mathbf{k}})$ appearing in Eq. (14) with a finite-width representation $\tilde{\delta}(\epsilon - \epsilon_{j\mathbf{k}})$. The LDOS used to train our ML model is then obtained by evaluating

$$D(\epsilon, \mathbf{r}; \mathbf{R}) = N_k^{-1} \sum_k \sum_j^{N_s} |\phi_{j\mathbf{k}_k}(\mathbf{r}; \mathbf{R})|^2 \tilde{\delta}(\epsilon - \epsilon_{j\mathbf{k}_k}) \quad (16)$$

on a finite grid of ϵ values. In our calculations, we use an evenly spaced ϵ grid with a spacing of 0.1 eV, a minimum value of -10 eV, and a maximum value of $+15$ eV. Hence, the data for each grid point consist of a vector of 250 scalar values. Kohn-Sham orbitals with eigenvalues significantly outside the energy range where we evaluate the LDOS make negligible contributions to the calculated LDOS values, and ideally we would include a sufficient number N_s of the lowest Kohn-Sham orbitals in Eqs. (14) and (16) to span this energy range. However, due to memory constraints on our $8 \times 8 \times 8$ k -point DFT calculations, we were only able to calculate $N_s = 576$ orbitals. At ambient density, this provides an accurate LDOS up to $\epsilon \approx 10$ eV, and the relevant Fermi levels are at $\mu \approx 7.7$ eV. The tails of $f^\beta(\epsilon)$ should be negligible above $\epsilon \approx 10$ eV for $T \lesssim 2500$ K. At higher electronic temperatures, it would become necessary to increase N_s when generating the training data in order to get an accurate LDOS over a wider energy range.

The challenge mentioned above arises in the choice of the δ function representation $\tilde{\delta}$ and the k -point sampling scheme $\mathbf{k}_1, \dots, \mathbf{k}_{N_k}$ for a given ϵ grid. If the function $\tilde{\delta}$ is too wide, it leads to systematic errors in the LDOS and derived quantities such as the density, band energy, and total energy. However, if $\tilde{\delta}$ is too narrow or too few k points are used to sample the BZ, then only a few values of $\epsilon_{j\mathbf{k}_k}$ make substantial contributions to the LDOS at each point on the ϵ grid, and there is a large amount of noise in the resulting LDOS.

Figure 2 demonstrates these trade-offs for a Gaussian representation $\tilde{\delta}(x) = \exp(-x^2/\sigma^2)/\sqrt{\pi\sigma^2}$ by plotting the DOS $\bar{D}[D](\epsilon; \mathbf{R})$ in the upper panel and the error in the band energy calculated from the DOS in the lower panel. The top panel shows that using fewer k points or a narrower $\tilde{\delta}$ leads to unphysical noise in the calculated DOS, while the bottom panel shows that using a wider $\tilde{\delta}$ leads to an increasing error in the band energy. The lower panel also shows that the noise in the DOS that arises with a narrower $\tilde{\delta}$ is reflected in an increasing variability in the band energy error.

Since a larger k -point sampling allows a narrower $\tilde{\delta}$ to be used without introducing excessive noise into the LDOS, we can minimize the errors in quantities calculated from the LDOS by using as many k points as possible. In the LDOS calculations used to train our ML model, we used $8 \times 8 \times 8$ k -point sampling, which was the largest k -point sampling that was computationally tractable for our 256-atom cells. We then chose a Gaussian $\tilde{\delta}$ with $\sigma = 0.2$ eV (indicated in the lower panel of Fig. 2), which gives a relatively small and very consistent error in derived quantities. For example, the average error in the band energies calculated from the resulting LDOS (relative to fully converged DFT) is -4.483 meV/atom when the average is taken over all of our snapshots. The average error becomes -4.466 , -4.500 , and -4.483 meV/atom when

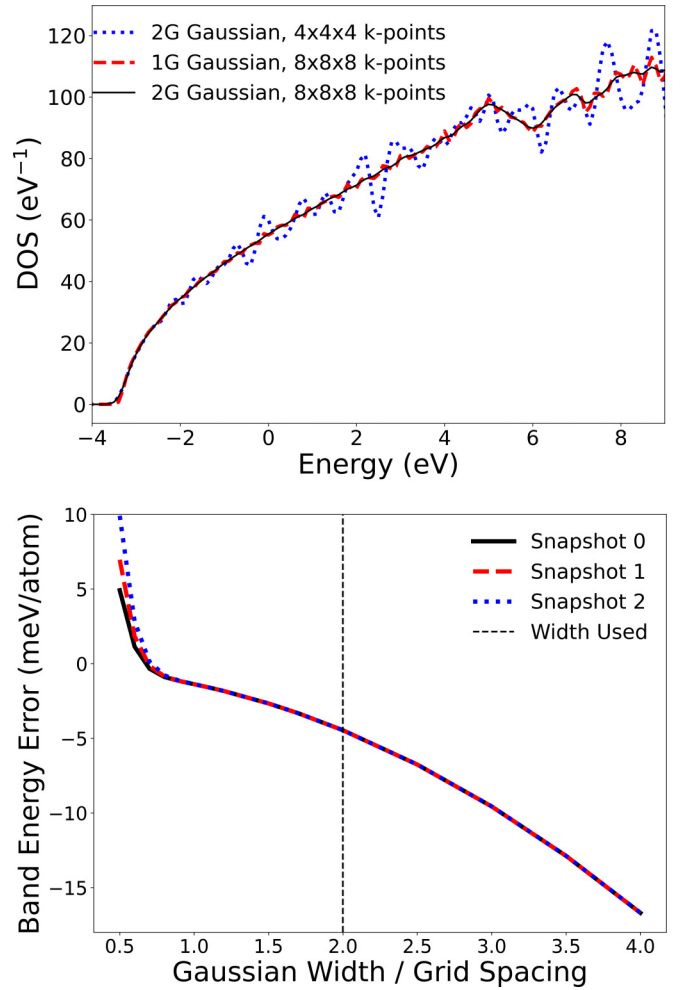


FIG. 2. The DOS calculated for one of the 256-atom aluminum snapshots at ambient temperature (298 K) using different Gaussian smearing widths and k -point sampling schemes (top) and the error in the band energy as a function of the Gaussian smearing width for three different snapshots (bottom). In the upper panel, “1G Gaussian” indicates results obtained from a Gaussian $\tilde{\delta}(x)$ with $\sigma = 0.1$ eV (equal to the ϵ grid spacing), while “2G Gaussian” indicates that $\sigma = 0.2$ eV (twice the ϵ grid spacing) was used. Results are shown for $4 \times 4 \times 4$ and $8 \times 8 \times 8$ Monkhorst-Pack k -point sampling in the upper panel, and $8 \times 8 \times 8$ sampling in the lower panel.

the average is restricted to the 298 K, 933 K solid, and 933 K liquid snapshots, respectively. These results show that the error changes very little between different temperatures and structures. These errors are also very consistent for different structures at the same temperature and phase. In particular, the root-mean-square variations in the errors are 0.004, 0.011, and 0.017 meV/atom for the 298 K, 933 K solid, and 933 K liquid snapshots, respectively. Our choices for the ϵ grid spacing and the width of the Gaussian broadening turn out to be similar to those made in Ref. [51].

In electronic structure calculations, consistent errors in calculated energies (for example, the errors due to incomplete convergence with respect to the plane-wave cutoffs in traditional DFT calculations) are generally not a problem. Instead, quantities of interest almost always involve the difference

between two energies, and thus it is the variation of the errors between different structures that affects actual results. With the k -point sampling and the choice of $\tilde{\delta}$ described above, the errors differ by only about one-thousandth of the room-temperature thermal energy, and such consistent errors are very unlikely to impact quantities of interest. Furthermore, these errors are independent of the errors introduced by using ML to approximate the LDOS. In particular, a perfect ML approximation would reproduce the results calculated from the DFT LDOS (defined in the same manner as in the generation of our training data) rather than the results calculated by exact DFT. Finally, we believe that further work will be able to identify methods that use DFT to calculate even more accurate LDOS training data with less computational effort than our $8 \times 8 \times 8$ k -point calculations. Some ideas include Fourier interpolation of the Kohn-Sham bands and orbitals, or an approach analogous to the tetrahedron method [83]. The remainder of this paper will focus on the effects of approximating the LDOS with ML. In light of the above considerations, errors will be quoted relative to results calculated from the DFT LDOS rather than exact DFT.

D. Evaluation of energies from the local density of states

A key step in calculating the total energy defined in Eq. (9) from the LDOS is the evaluation of several integrals in Eqs. (10)–(12) that have the form

$$I = \int d\epsilon g(\epsilon) D(\epsilon), \quad (17)$$

for some function $g(\epsilon)$. The evaluation of Eq. (17) poses a challenge, because $g(\epsilon)$ changes rapidly compared to $D(\epsilon)$ in several cases of interest. We provide a solution in terms of an analytical integration, as illustrated in Appendix A. Given our analytical integration technique, it is straightforward to evaluate the total energy using Eq. (9) and the standard methods of DFT in an accurate and numerically stable manner.

E. Fingerprint generation

We assume that the LDOS at any point in space can be approximated by a function that depends only on the positions and chemical identities of atoms within some finite neighborhood of the point. In order to approximate this function using ML, we need to construct a fingerprint that maps the neighborhood of any point to a set of scalar values called *descriptors*. On physical grounds, good descriptors must satisfy certain minimum requirements: (i) invariance under permutation, translation, and rotation of the atoms in the neighborhood and (ii) continuous differentiable mapping from atomic positions to descriptors, especially at the boundary of the neighborhood. While these requirements exclude some otherwise appealing choices, such as the Coulomb matrix [42], the space of physically valid descriptors is still vast. In the context of ML-IAPs, the construction of good atomic neighborhood descriptors has been the subject of intense research. Recent work by Drautz [84] has shown that many prominent descriptors [28,32,72] belong to a larger family that are obtained from successively higher-order terms in an expansion of the local atomic density in cluster integrals. The SNAP bispectrum descriptors [32]

that we use in this work correspond to clusters of three neighbor atoms yielding four-body descriptors.

In contrast to prior work on ML-IAPs in which descriptors are evaluated on atom-centered neighborhoods, here we evaluate the SNAP descriptors on the same $200 \times 200 \times 200$ Cartesian grid points at which we evaluated the LDOS training data [see Eq. (16) above]. For each grid point, we position it at the origin and define local atom positions in the neighborhood relative to it. The total density of neighbor atoms is represented as a sum of δ functions in a three-dimensional space:

$$\rho(\mathbf{r}) = \delta(\mathbf{0}) + \sum_{r_k < R_{\text{cut}}^{v_k}} f_c(r_k; R_{\text{cut}}^{v_k}) w_{v_k} \delta(\mathbf{r}_k), \quad (18)$$

where \mathbf{r}_k is the position of the neighbor atom k of element v_k relative to the grid point. The w_v coefficients are dimensionless weights that are chosen to distinguish atoms of different chemical elements v , while a weight of unity is assigned to the location of the LDOS point at the origin. The sum is over all atoms k within some cutoff distance $R_{\text{cut}}^{v_k}$ that depends on the chemical identity of the neighbor atom. The switching function $f_c(r; R_{\text{cut}}^{v_k})$ ensures that the contribution of each neighbor atom goes smoothly to zero at $R_{\text{cut}}^{v_k}$. Following Bartók *et al.* [28], the radial distance r_k is mapped to a third polar angle θ_0 defined by

$$\theta_0 = \theta_0^{\text{max}} \frac{r_k}{R_{\text{cut}}^{v_k}}. \quad (19)$$

The additional angle θ_0 allows the set of points \mathbf{r}_k in the three-dimensional (3D) ball of possible neighbor positions to be mapped onto the set of points (θ, ϕ, θ_0) on the unit 3-sphere. The neighbor density function can be expanded in the basis of four-dimensional (4D) hyperspherical harmonic functions \mathbf{U}^j ,

$$\rho(\mathbf{r}) = \sum_{j=0, \frac{1}{2}, \dots}^{\infty} \mathbf{u}_j \cdot \mathbf{U}_j(\theta_0, \theta, \phi), \quad (20)$$

where \mathbf{u}_j and \mathbf{U}_j are rank $(2j+1)$ complex square matrices, \mathbf{u}_j are Fourier expansion coefficients given by the inner product of the neighbor density with the basis functions \mathbf{U}_j of degree j , and the \cdot symbol indicates the scalar product of the two matrices. Because the neighbor density is a weighted sum of δ functions, each expansion coefficient can be written as a sum over discrete values of the corresponding basis function

$$\mathbf{u}_j = \mathbf{U}_j(\mathbf{0}) + \sum_{r_k < R_{\text{cut}}^{v_k}} f_c(r_k; R_{\text{cut}}^{v_k}) w_{v_k} \mathbf{U}_j(\theta_0, \theta, \phi). \quad (21)$$

The expansion coefficients \mathbf{u}_j are complex valued and they are not directly useful as descriptors because they are not invariant under rotation of the polar coordinate frame. However, scalar triple products of the expansion coefficients

$$\mathbf{B}_{j_1 j_2 j} = \mathbf{u}_{j_1} \otimes_{j_1 j_2 j} \mathbf{u}_{j_2} \cdot (\mathbf{u}_j)^* \quad (22)$$

are real valued and invariant under rotation [28]. The symbol $\otimes_{j_1 j_2 j}$ indicates a Clebsch-Gordan product of matrices of rank j_1 and j_2 that produces a matrix of rank j , as defined in our original formulation of SNAP [32]. These invariants are the components of the bispectrum. They characterize the strength

of density correlations at three points on the 3-sphere. The lowest-order components describe the coarsest features of the density function, while higher-order components reflect finer detail. The bispectrum components defined here have been shown to be closely related to the four-body basis functions of the atomic cluster expansion introduced by Drautz [84].

For computational convenience, the LAMMPS implementation of the SNAP descriptors on the grid includes all unique bispectrum components $B_{j_1 j_2 j}$ with indices no greater than J_{\max} . For the current study we chose $J_{\max} = 5$, yielding a fingerprint vector consisting of 91 scalar values for each grid point. The cutoff distance and element weight for the aluminum atoms were set to 4.676 Å and 1.0, respectively.

F. Machine-learning workflow and neural network model

Neural network models, especially deep neural networks, are profoundly successful in numerous tasks, such as parameter estimation [85,86], image classification [87,88], and natural language processing [89,90]. Moreover, neural networks may act as function approximators [91,92]. A typical *feed-forward* neural network is constructed as a sequence of transformations, or *layers*, with the form

$$v_{(n,1)}^{\ell+1} = \varphi(\mathbf{W}_{(n,m)}^{\ell} v_{(m,1)}^{\ell} + b_{(n,1)}^{\ell}), \quad (23)$$

where $v_{(m,1)}^{\ell}$ is the intermediate column vector of length m at layer ℓ , $\mathbf{W}_{(n,m)}^{\ell}$ is a matrix of size $n \times m$, $b_{(n,1)}^{\ell}$ is a bias vector of length n , and $v_{(n,1)}^{\ell+1}$ is the next column vector of length n in the sequence. The first and the last layers accept *input* vectors, v^0 , and produce *output* predictions, v^L , respectively. There can be any arbitrary number $L - 1$ of *hidden* layers constructed between the input layer and the output layer. After each layer, a nonlinear *activation* unit φ transforms its input. The composition of these layers is expected to learn complex, nonlinear mappings.

Given a data set of N_s input vectors and target output vectors, (v^0, \bar{v}^L) , we seek to find the optimal W^{ℓ} and b^{ℓ} , $\forall \ell \in [0, L - 1]$, such that the neural network minimizes the root-mean-squared loss between the neural network predictions v^L and targets \bar{v}^L . In this work, the network predicts LDOS at each grid point using the SNAP fingerprint as input. The specific feed-forward networks we consider have input fingerprint vectors of length 91 scalar values and output LDOS vectors of length 250 scalar values. The grid points are defined on a uniform $200 \times 200 \times 200$ Cartesian grid, or 8 million total points per snapshot.

Neural network models are created in two stages, *training* and *validation*, using separate training and validation data sets. The training phase performs gradient-based updates, or *backpropagation* [92], to W^{ℓ} and b^{ℓ} using a small subset, or *minibatch*, of column vectors in the training data set. Steps taken in the direction of the gradient are weighted by a user-provided *learning rate* (LR). Successive updates incrementally optimize the neural network. An *epoch* is one full pass through the training data set for the training phase. After each epoch, we perform the validation phase, which is an evaluation of the current model's accuracy on a validation data set. The training process continues alternating between training and validation phases until an *early stopping* termination

criterion [93] is satisfied. Early stopping terminates the learning process when the root-mean-squared loss on the validation data set does not decrease for a successive number of epochs, or *patience* count. Increasing validation loss indicates that the model is starting to over-fit the training data set and will have worse generalization performance if the training continues.

Next, the hyperparameters, which determine the neural network architecture and optimizer specifications, must be optimized to maximize model accuracy (and minimize loss). Due to the limited throughput, we make the assumption that all hyperparameters are independent. Independence allows the hyperparameter tuning to optimize one parameter at a time, greatly reducing the total rounds of training. We then ensure local optimality of the hyperparameter minimizer by direct search [94].

The ML-DFT model uses as input vectors the SNAP fingerprints defined at real-space grid points. The target output vectors of the ML-DFT model are the QUANTUM ESPRESSO-generated LDOS vectors defined at the corresponding real-space grid points. Importantly, each fingerprint and LDOS vector is treated as independent of any other grid point. Further, the training data may include multiple aluminum snapshots, in which case the number of grid points is equal to the number of snapshots times the number of grid points per snapshot. No distinction is made in the model by grid points belonging to separate snapshots. The validation snapshot is always a full set of grid points from a single snapshot.

For inference, all remaining snapshots are used to test the generalization of the ML-DFT model. Consider a single snapshot where we wish to predict the total energy. At each grid point, a SNAP fingerprint vector is used to predict an LDOS vector. Using the steps in Sec. II D and Appendix A, the full set of ML-predicted LDOS vectors produce a total energy.

In summary, the ML-DFT training workflow is provided in Table I. The ML-DFT models are built on top of the PyTorch [95] framework and the NumPy [96] and Horovod [97] packages. PyTorch provides the necessary infrastructure for neural network construction and efficiently running the training calculations on GPU machines. The NumPy package contains many performant tools for manipulating tensor data, and the Horovod package enables data-parallel training for MPI-based scalability.

III. RESULTS

The difficulty of predicting the electronic structure of materials is strongly dependent on the diversity of local atomic configurations considered in training and testing. During the development stage of this work we focused on the relatively simple case of solid aluminum at ambient temperature and density (298 K, 2.699 g/cm³). We then turned our attention to the more challenging case of solid and liquid aluminum near the melting point (933 K, 2.699 g/cm³) and we present results for that case here.

Specifically, we train and tune a collection of unique ML-DFT models for each temperature. In the 298 K case, we train a single model using one snapshot, validate on one snapshot, and test on eight snapshots. At 933 K, we train eight models with different combinations of liquid and solid snapshots as

TABLE I. ML-DFT model training workflow.

(1)	Training data generation
(i)	Run MD simulation to obtain atomic configurations at a specific temperature and density.
(ii)	Run QUANTUM ESPRESSO DFT calculation on the atomic configurations.
(iii)	Perform QUANTUM ESPRESSO postprocessing of the Kohn-Sham orbitals to obtain the LDOS on the Cartesian grid.
(iv)	Generate SNAP fingerprints on a uniform Cartesian grid from the atomic configurations.
(2)	ML-DFT model training
(i)	Training and validation set, ML model, and initial hyperparameter selection.
(ii)	Normalization and standardization of the training data.
(iii)	Train ML-DFT model to convergence.
(iv)	Adjust hyperparameters for increased accuracy and repeat training for a new model.
(3)	ML-DFT analysis
(i)	Evaluate unseen test snapshots on trained and tuned ML-DFT models.
(ii)	Calculate quantities of interest from the LDOS predicted by ML-DFT.

the training set. The first three cases use four training snapshots with either four liquid snapshots, four solid snapshots, or two liquid and two solid snapshots. The second three cases, similarly, use eight training snapshots with either eight liquid, eight solid, or four liquid and four solid snapshots. Finally, there are two models that use 8 or 12 training snapshots and a validation set that includes both liquid and solid snapshots. Each 933 K model created is then tested on the remaining 933 K snapshots. In order to study the variability in LDOS between the three groups (298 K solid, 933 K solid, and 933 K liquid), we reduce the dimensionality of the LDOS data sets and study them using principal component analysis (PCA).

Further technical details on fingerprint generation, neural network architecture, ML training, optimization of hyperparameters, and the variability in LDOS outputs are given in Appendix B. We provide results for the ambient case in Appendix C.

A. Local density of states predictions for solid and liquid aluminum from ML-DFT

First we assess the accuracy of ML-DFT for spatially resolved and energy-resolved quantities. To this end, Fig. 3 illustrates the DOS and its errors in the relevant energy range

from -5 to 10 eV, where the Fermi energy is at 7.689 eV in the liquid and 7.750 eV in the solid phase. The DOS prediction is computed from Eq. (13) using the ML-DFT predicted LDOS. In both liquid (left) and solid (right) phases, the ML-DFT model reproduces the DOS to very high accuracy. The illustration of the errors (bottom panels) confirms the ML-DFT model's accuracy and shows the absence of any unwanted error cancellation over the energy grid. Figure 3 also shows that the Van Hove singularities [82] at and above ≈ 5 eV, which show up strongly in the 298 K DOS in Fig. 2, are noticeably smeared out by thermal fluctuations in the atomic positions in the 933 K solid-phase DOS and essentially gone in the 933 K liquid. Figure 4 shows the DFT target electronic density versus the predicted electronic density for both liquid (left) and solid (right) phases at each point on the Cartesian grid. The frequency of the predicted results is indicated by the marker color. The ML-DFT prediction is computed from Eq. (10) using the ML-DFT predicted LDOS. The alignments along the diagonal exhibit small standard deviations for the liquid and solid test snapshot. It illustrates how well the ML-DFT model reproduces the electronic density in a systematic manner with high accuracy, despite the fact that the electronic structure is qualitatively very distinct in the solid and liquid phases of aluminum. Furthermore,

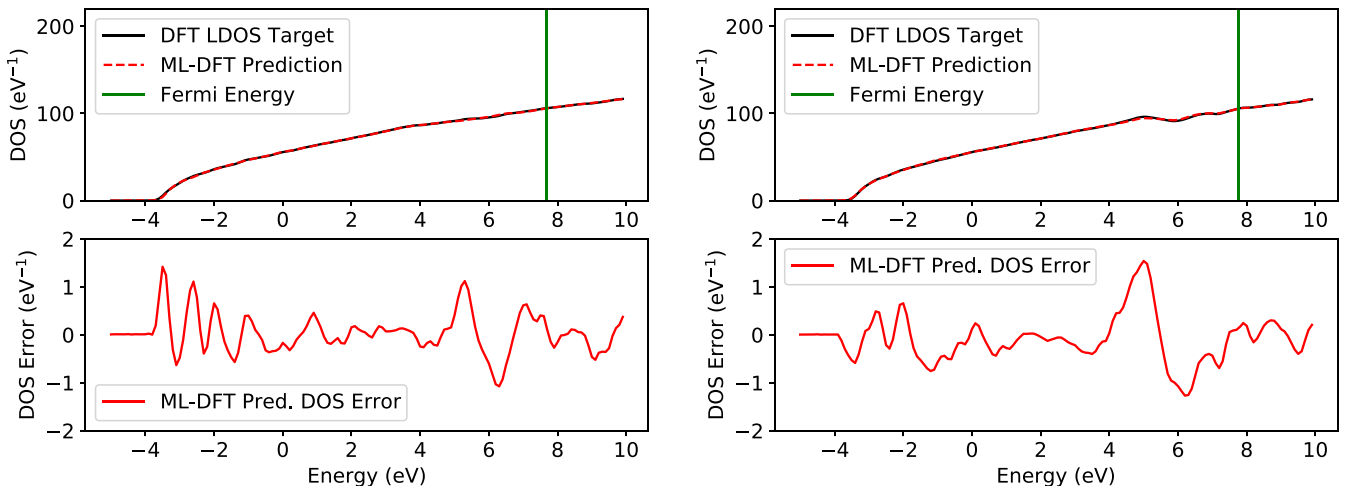


FIG. 3. Liquid (left) and solid (right) test snapshot DOS comparisons for the ML-DFT model trained on six liquid and six solid snapshots at 933 K. The vertical lines indicate the Fermi energy which is 7.689 eV for the liquid snapshot and 7.750 eV for the solid snapshot.




FIG. 4. Liquid (left) and solid (right) test snapshot electron density comparisons for the ML-DFT model trained on six liquid and six solid snapshots at 933 K. The maximum absolute errors for the liquid and solid test snapshot are $0.00601 e^{-}/\text{\AA}^3$ and $0.00254 e^{-}/\text{\AA}^3$, respectively; mean absolute errors are $0.000284 e^{-}/\text{\AA}^3$ and $0.000167 e^{-}/\text{\AA}^3$, respectively; and mean absolute percentage errors are 1.12% and 0.66%, respectively. Marker color indicates the frequency.

maximum absolute errors for the liquid and solid test snapshot are $0.00601 e^{-}/\text{\AA}^3$ and $0.00254 e^{-}/\text{\AA}^3$, respectively. The mean absolute errors are an order of magnitude lower at $0.000284 e^{-}/\text{\AA}^3$ and $0.000167 e^{-}/\text{\AA}^3$, respectively. The mean absolute percentage errors are 1.12% and 0.66%, respectively.

B. Single-phase solid and liquid aluminum ML-DFT models

We demonstrate the errors in total energy and band energy on models trained on data from the same phase, either solid or liquid. Figure 5 and Table II show these results. When training, validation, and test sets are all from the same phase, either solid or liquid, the overall ML-DFT accuracy is very high. The mean absolute error in the total energy for the eight-snapshot liquid model is 9.31 meV/atom. Moreover, the eight-snapshot solid model has a mean absolute error of 15.71 meV/atom. In both cases, the main contribution to the error is due to the band energy error which is 5.21 meV/atom in the

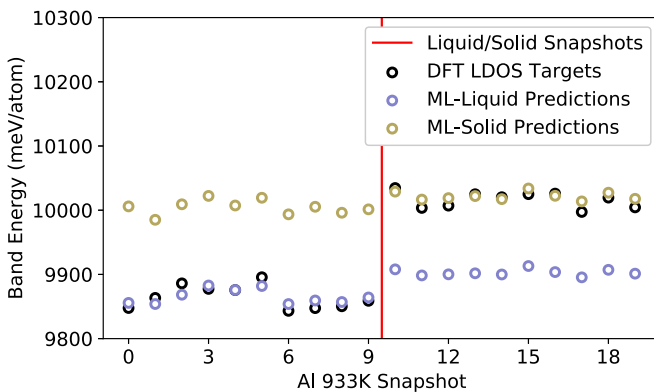


FIG. 5. Band energy comparisons for ML-DFT models trained on either eight liquid or eight solid snapshots at 933 K. The red line divides the ten liquid (left) and ten solid (right) snapshots considered. The ML-Liquid training snapshots are snapshots 0–7 and the validation snapshot is snapshot 8. The ML-Solid training snapshots are snapshots 10–17 and the validation snapshot is snapshot 18.

liquid test set and 13.37 meV/atom in the solid. Both models have a similar mean absolute percentage error, 0.01% in the liquid model and 0.03% in the solid model.

However, training on a single phase results in poor predictive accuracy for the phase unseen in training, as is shown clearly in Fig. 5 and Table II. For the ML-DFT model trained on eight liquid snapshots, the mean absolute total energy error in the solid test set rises to 111.41 meV/atom. For the model trained on eight solid snapshots, the mean absolute total energy error in the liquid test set rises to 123.29 meV/atom. Again, the major contribution to the error stems from the band energy which is 113.40 meV/atom in the solid test set and 139.96 meV/atom in the liquid.

C. Hybrid liquid-solid aluminum ML-DFT models

In Fig. 6 and Table II we demonstrate our central result, a single ML-DFT model that is capable of yielding accurate results for both liquid and solid aluminum at a temperature of 933 K with accuracy sufficient for meaningful atomistic simulations of electronic, thermodynamic, and kinetic behavior of materials. Using the hyperparameters listed in Table III, the most accurate dual-phase *hybrid* ML-DFT model was trained using six liquid and six solid snapshots (snapshots 0–5 and 10–15) as the training data set and one solid and one liquid snapshot (snapshots 6 and 16) as a validation set. This ML-DFT model achieves state-of-the-art ML accuracy for all test snapshots considered.

Overall, Table II summarizes the total energy [Eq. (9)] and band energy [Eq. (11)] errors for each of the eight ML-DFT models. Training and validation sets determine the specific ML-DFT model and the test sets are used to measure the model's ability to generalize to liquid or solid phase snapshots.

Most notably, our overall most accurate hybrid ML-DFT model predicts the total energy over both solid and liquid phases with a mean absolute error of 13.04 meV/atom in the solid test set and 31.60 meV/atom in the liquid test set, as illustrated in Fig. 6 and listed in bold in the last two rows of

TABLE II. Band energy and total energy test errors from different ML-DFT models inferred for aluminum snapshots at 933 K. The training set consists of snapshots in liquid or solid states. When trained, validated, and tested on the same phase, either liquid (highlighted in blue) or solid (highlighted in red), the overall ML-DFT accuracy is very high. The training on a single phase results in poor predictive accuracy for the phase unseen in training, as expected. In contrast, our main result, the hybrid ML-DFT model, predicts the total energy over both solid and liquid phases with accuracy sufficient for meaningful atomistic simulations (highlighted in green). Maximum Absolute errors (Max. AE), mean absolute error (MAE) and mean absolute percentage error (MAPE) are shown.

Training set (Validation set)	Test set	Band energy			Total energy		
		Max. AE (meV/atom)	MAE (meV/atom)	MAPE (%)	Max. AE (meV/atom)	MAE (meV/atom)	MAPE (%)
4 liquid	5 liquid	13.57	9.80	0.0994	12.39	7.13	0.0113
(1 liquid)	10 solid	113.50	100.84	1.0067	106.62	91.54	0.1452
8 liquid	1 liquid	5.21	5.21	0.0529	9.31	9.31	0.0148
(1 liquid)	10 solid	126.69	113.40	1.1320	125.07	111.41	0.1768
4 solid	10 liquid	142.63	126.99	1.2875	142.95	126.78	0.2015
(1 solid)	5 solid	13.15	8.44	0.0843	14.45	9.81	0.0156
8 solid	10 liquid	158.25	139.96	1.4190	153.45	123.29	0.1959
(1 solid)	1 solid	13.37	13.37	0.1336	15.71	15.71	0.0249
2 liquid + 2 solid	8 liquid	73.76	62.14	0.6299	57.34	48.47	0.0770
(1 solid)	7 solid	37.26	28.22	0.2816	41.53	29.50	0.0468
4 liquid + 4 solid	6 liquid	67.28	58.23	0.5906	61.29	51.37	0.0816
(1 solid)	5 solid	25.33	16.16	0.1613	26.98	15.11	0.0240
4 liquid + 4 solid	5 liquid	36.84	28.54	0.2896	28.39	21.33	0.0339
(1 liquid + 1 solid)	5 solid	44.94	37.19	0.3711	47.30	38.57	0.0612
6 liquid + 6 solid	3 liquid	21.34	17.11	0.1737	15.76	13.04	0.0207
(1 liquid + 1 solid)	3 solid	39.33	33.56	0.3354	39.23	31.60	0.0501

Table II. In comparison with the single-phase ML-DFT models in Fig. 5, the hybrid ML-DFT model is able to generalize to both liquid and solid snapshots. This is adequate to resolve the solid-liquid total energy difference (110 meV/atom) and is approaching the 5 meV/atom accuracy of the best ML-IAPs used in large-scale MD simulations [98].

Furthermore, the sensitivity of ML-DFT towards the volume of training data is also assessed in Table II. Reducing the training data set to only two liquid and two solid snapshots degrades the accuracy for both the band and total energies as expected. While the mean absolute error of 29.50 meV/atom in the solid phase is about the same as in the larger training set, the error in the liquid phase rises to 48.47 meV/atom. This emphasizes the fact that a desired accuracy in the ML-DFT

model can be approached systematically by increasing the volume of training data.

Interestingly, in Fig. 6, the training snapshots do not necessarily have the lowest error as might be expected. The band and total energy errors for training snapshot 0, for example, are slightly larger than the errors for test snapshots 8 and 9. For the solid snapshots, both training snapshots (10–15) and test snapshots (17–19) have very similar errors in both the band energy and total energy. Similar observations can be made for Fig. 5. A greater number of liquid and solid test snapshots will only further refine the statistics of the ML-DFT model accuracy. Future investigations will include both sensitivity analysis and thorough validation of the ML-DFT model, particularly in the context of large-scale MD simulations.

TABLE III. Optimized hyperparameters for the 298 K and 933 K models.

Model parameter	298 K model	933 K models
Fingerprint length	91	91
Fingerprint scaling	Element-wise standardization	Element-wise standardization
LDOS length	250	250
LDOS scaling	Max normalization	Max normalization
Optimizer	Adam	Adam
Minibatch size	1000	1000
Learning rate (LR)	1×10^{-5}	5×10^{-5}
LR schedule	4 epochs	4 epochs
Early stopping	8 epochs	8 epochs
Max layer width	800	4000
Layers	5	5
Activation function	LeakyReLU	LeakyReLU
Total weights	1.62×10^6	3.34×10^7

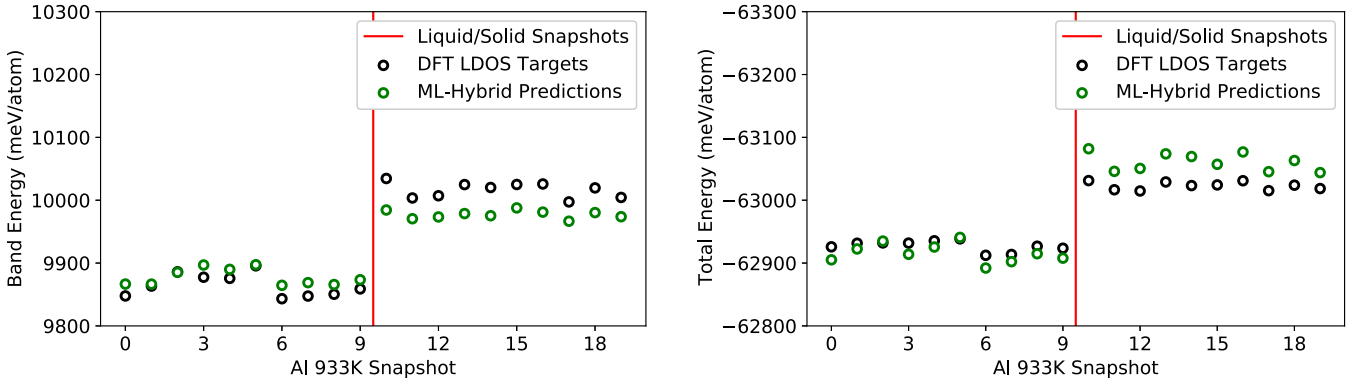


FIG. 6. Band energy (left) and total energy (right) comparisons for the ML-DFT model trained on six liquid and six solid snapshots at 933 K. The red line in each figure divides the ten liquid (left) and ten solid (right) snapshots considered. The ML-Hybrid training snapshots are snapshots 0–5 and 10–15 and the validation snapshots are snapshots 6 and 16.

IV. DISCUSSION

In the present work, we establish an ML framework that eliminates the most computationally demanding and poorly scaling components of DFT calculations. Once trained with DFT results for similar systems, it can predict many of the key results of a DFT calculation, including the LDOS, DOS, density, band energy, and total energy, with good accuracy. After training, the only input required by the model is the atomic configuration, i.e., the positions and chemical identities of atoms. We have been able to train a single model that accurately represents both solid and liquid properties of aluminum. The ML-DFT model is distinct from previous ML methodologies, as it is centered on materials modeling in support of accurate MD-based predictions for the evolution of both electronic and atomic properties. We represent materials descriptors (such as the LDOS) on a regular grid as opposed to basis functions [49] and use a specific formulation in terms of the LDOS to reconstruct the total energy. Our approach also generalizes to nonzero electronic temperatures in a straightforward way, which allows it to be easily applied to materials at extreme conditions.

Furthermore, the very recently published work [52] stopped short of demonstrating that the ML-predicted LDOS can be utilized to accurately calculate energies and forces needed to enable MD calculations. Further work on computing accurate forces of materials under elevated temperatures and pressures will follow, as well as a detailed exploration of ML-DFT’s accuracy when scaled to $O(10^4)$ atoms, well beyond the reach of standard DFT.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the

U.S. Government. This work was partly funded by the Center for Advanced Systems Understanding (CASUS) which is financed by the German Federal Ministry of Education and Research (BMBF) and by the Saxon State Ministry for Science, Art, and Tourism (SMWK) with tax funds on the basis of the budget approved by the Saxon State Parliament. We thank Warren Davis and Joshua Rackers for helpful conversations and Mitchell Wood for visualizing the SNAP fingerprint grid.

APPENDIX A: ANALYTICAL EVALUATION OF DERIVED QUANTITIES FROM THE LOCAL DENSITY OF STATES

Consider the numerically challenging integral in Eq. (17). When $g = f^\beta(\epsilon)$, this integral gives the electronic density and the total number of electrons, when $g = \epsilon f^\beta(\epsilon)$ the result is the band energy, and when $g = g^S \equiv \beta^{-1} \{f^\beta(\epsilon) \log[f^\beta(\epsilon)] + [1 - f^\beta(\epsilon)] \log[1 - f^\beta(\epsilon)]\}$ we obtain the entropy contribution to the energy. When $D(\epsilon) = \bar{D}[D](\epsilon; \mathbf{R})$, Eq. (17) results in a scalar quantity, and we can evaluate the total number of electrons, the band energy, or the entropy contribution to the total energy. Alternatively, when $D(\epsilon) = D(\epsilon, \mathbf{r}; \mathbf{R})$, Eq. (17) results in a field that depends on the position \mathbf{r} with the electronic density being the most important result.

Our ML model gives us $D(\epsilon_i)$ evaluated on a grid of energy values ϵ_i . We will extend $D(\epsilon_i)$ to a $D(\epsilon)$ defined for all ϵ between ϵ_0 and ϵ_N by linear interpolation. This allows us to perform the required integrals analytically. It is necessary to treat these integrals carefully because $f^\beta(\epsilon)$ changes rapidly compared to $D(\epsilon)$ for many systems of interest. In particular, $f^\beta(\epsilon)$ changes from one to zero over a few $k_B T$, while the spacing of the ϵ_i grid is several times $k_B T$ at room temperature. In order to evaluate these integrals analytically, we can use the representation of $f^\beta(\epsilon)$ and $g^S(\epsilon)$ in terms of the polylogarithm special functions $\text{Li}_n(x)$,

$$f^\beta(x) = 1 + \text{Li}_0(-e^x), \quad (\text{A1})$$

$$g^S(x) = \beta^{-1} [-x \text{Li}_0(-e^x) + \text{Li}_1(-e^x)], \quad (\text{A2})$$

where $x = \beta(\epsilon - \mu)$. It is useful to define a series of integrals giving moments of f^β and g^S with respect to $\epsilon - \mu$ and evaluate these integrals using integration by parts and the properties




FIG. 7. Variance of the first and second principal component scores for each group of snapshots, 298 K solids, 933 K solids, and 933 K liquids. Each group contains ten snapshots. The boxes enclose the second and third quartiles, and the whiskers indicate the minimum and maximum values.

of $\text{Li}_n(x)$. In particular, the relationship

$$\frac{d \text{Li}_n(-e^x)}{dx} = \text{Li}_{n-1}(-e^x) \quad (\text{A3})$$

is useful. The required integrals are

$$F_0 \equiv \beta^{-1} \int dx f^\beta(x) = \beta^{-1} [x + \text{Li}_1(-e^x)], \quad (\text{A4})$$

$$F_1 \equiv \beta^{-2} \int dx x f(x) = \beta^{-2} \left[\frac{x^2}{2} + x \text{Li}_1(-e^x) - \text{Li}_2(-e^x) \right], \quad (\text{A5})$$

$$F_2 \equiv \beta^{-3} \int dx x^2 f(x) = \beta^{-3} \left[\frac{x^3}{3} + x^2 \text{Li}_1(-e^x) - 2x \text{Li}_2(-e^x) + 2 \text{Li}_3(-e^x) \right], \quad (\text{A6})$$

$$S_0 \equiv \beta^{-1} \int dx g^S(x) = \beta^{-2} [-x \text{Li}_1(-e^x) + 2 \text{Li}_2(-e^x)], \quad (\text{A7})$$

$$S_1 \equiv \beta^{-2} \int dx x g^S(x) = \beta^{-3} [-x^2 \text{Li}_1(-e^x) + 3x \text{Li}_2(-e^x) - 3 \text{Li}_3(-e^x)]. \quad (\text{A8})$$

Quantities of interest with the form of Eq. (17) can then be evaluated as a weighted sum over the DOS or LDOS evaluated at the energy grid points

$$I = \sum_i w_i D(\epsilon_i), \quad (\text{A9})$$

where the weights w_i are given by

$$w_i = [I_0(\epsilon_{i+1}) - I_0(\epsilon_i)] \left[1 + \frac{\epsilon_i - \mu}{\epsilon_{i+1} - \epsilon_i} \right] + [I_0(\epsilon_i) - I_0(\epsilon_{i-1})] \left[1 - \frac{\epsilon_i - \mu}{\epsilon_i - \epsilon_{i-1}} \right] - \frac{I_1(\epsilon_{i+1}) - I_1(\epsilon_i)}{\epsilon_{i+1} - \epsilon_i} + \frac{I_1(\epsilon_i) - I_1(\epsilon_{i-1})}{\epsilon_i - \epsilon_{i-1}}. \quad (\text{A10})$$

When $D(\epsilon) = \bar{D}[D](\epsilon; \mathbf{R})$, $I_0 = F_0$, and $I_1 = F_1$, we obtain $I = N_e$, the total number of electrons in the system. The Fermi level μ is then adjusted until N_e is equal to the total number of valence electrons in the system (i.e., the system is charge neutral). When $I_0 = F_1$ and $I_1 = F_2$, we obtain $I = E_b[D] - \mu N_e$, and we can easily obtain the band energy $E_b[D]$ by adding μN_e . When $I_0 = S_0$ and $I_1 = S_1$, we obtain $I = -\beta^{-1} S_S[D]$, the entropy contribution to the energy. Finally, when $D(\epsilon) = D(\epsilon, r; \mathbf{R})$, $I_0 = F_0$, and $I_1 = F_1$, we obtain $n[D](\mathbf{r})$, the electronic density. Given these quantities, it is straightforward to evaluate the total (free) energy using Eq. (9) and the standard methods of DFT.

TABLE IV. Timings in different steps contributing to the total evaluation of KS-DFT quantities based on neural network inference. The reported times are averaged over the inference of ten snapshots. The focus of this work which is to infer LDOS using the ML inference is shown in bold.

Step	Time (s), 298 K model	Time (s), 933 K model	Optimization notes
Generate fingerprint vectors from atomic positions	2909.75	2908.12	LAMMPS used only in serial mode, speedup possible by using parallelization
Load modules	0.36	0.36	
Infer LDOS from neural network	18.72	53.94	Only a single GPU was used
Integrate LDOS to DOS	1.57	1.44	
Determine Fermi level	2.23	2.40	
Integrate LDOS to density	1282.39	1277.66	Integration unoptimized and in serial
Evaluate total energy contributions from DOS	0.96	0.99	
Evaluate total energy contributions from density	87.50	87.17	QUANTUM ESPRESSO used only in serial mode, speedup possible by using parallelization

TABLE V. Band energy and total energy test errors from the ML-DFT models inferred for aluminum snapshots at 298 K. Maximum Absolute errors (Max. AE), and mean absolute error (MAE) are shown.

Training set (validation set)	Test set	Band energy		Total energy	
		Max. AE (meV/atom)	MAE (meV/atom)	Max. AE (meV/atom)	MAE (meV/atom)
1 solid (1 solid)	8 solid	3.82	2.48	5.60	3.07

APPENDIX B: MACHINE-LEARNING MODEL DETAILS

The 298 K systems include ten snapshots of SNAP fingerprint and QUANTUM ESPRESSO LDOS data. At 933 K, there are ten snapshots of liquid aluminum and ten snapshots of solid aluminum. For the fingerprints and LDOS data, each snapshot uses the same uniform Cartesian real-space grid ($200 \times 200 \times 200 = 8 \times 10^6$ points). Each snapshot, including both SNAP fingerprint and LDOS data, is 20.7 gigabytes.

We observe that the 933 K ML-DFT models require a greater amount of training data. Furthermore, the optimized 298 K model requires approximately 1.6 million weights, whereas the 933 K models require 33.4 million weights. The larger training set and greater model complexity can be attributed to greater LDOS variability at the higher temperature.

To facilitate comparing LDOS variability between the three groups (298 K solid, 933 K solid, and 933 K liquid), the dimensionality of the LDOS data sets was reduced using PCA. Because of the large amount of data ($30 \times 8 \times 10^6$ samples by 250 energy levels), PCA was performed incrementally using the scikit-learn PYTHON package [99]. The resulting first principal component explained approximately 81.7% of the variance in the LDOS, and the second explained an additional 7.9%.

The variance of the first two principal component scores was calculated for each of the 30 snapshots. Figure 7 shows these variances grouped by snapshot temperature and phase. The 933 K liquid snapshots are seen to possess the highest variance, followed by the 933 K solid snapshots. The 298 K solid snapshots have the lowest variance.

We perform a sequence of hyperparameter tunings once for each temperature using the optimization strategy described in Sec. II F. The final 933 K models are trained using the same hyperparameters. Further accuracy may be obtainable once training throughput is increased and multiple, independent hyperparameter optimizations are made possible. The optimized ML-DFT model parameters are given in Table III. Scaling of the input and output data prior to training is critical to ML-DFT model accuracy. We perform an *element-wise* standardization to mean zero and standard deviation 1 for the fingerprint inputs. For example, we standardize the first scalar of all fingerprint vectors in the training set, then we standardize the second scalar of all fingerprints, and so on. For the LDOS outputs, we perform a normalization using the maximum LDOS value of the entire training data set.

The hyperparameter convergence for the 298 K and 933 K models required 48 and 23 training iterations, respectively. Once the 298 K case was optimized, using those optimized hyperparameters as the 933 K optimization's initial iterate greatly accelerated convergence resulting in fewer training iterations. Within an iteration, the 298 K model requires 93 epochs. For the eight-snapshot 933 K cases, the single-phase models, either liquid or solid, required just 39–50 epochs to converge. Comparatively, the final hybrid liquid-solid model required 101 epochs to converge.

Due to the large quantity of training data in the 933 K case (12 training + 2 validation snapshot \times 20.7 gigabytes), the required training strategy for single-node machines is to *lazily*




FIG. 8. Band energy (left) and total energy (right) comparisons for the ML-DFT model trained on one solid snapshot at 298 K. The training and validation snapshot are snapshot 0 and 1, respectively.




FIG. 9. DOS comparison for a test snapshot using the ML-DFT model trained on snapshots at 298 K. The vertical line indicates the Fermi energy which is 7.797 eV.

load the snapshots into memory each epoch. This strategy is needed for machines where the CPU memory is insufficient and allows ML-DFT to train on any arbitrary number of snapshots. All ML-DFT models have been trained on a single node with one NVIDIA Tesla GPU. The 298 K network, with only one training and one validation snapshot, does not need to lazily load snapshots, and so each epoch requires approximately 151 s. For the 933 K networks, the total runtime of a single epoch with 12 training snapshots is approximately 76 min. With a sufficient number of nodes, we are able to stage each training snapshot on an independent node to alleviate the lazy loading data movement penalty completely. Future performance improvements to the workflow will target this bottleneck chiefly.

Finally, ML-DFT inference requires only 54 s to obtain a full grid of LDOS predictions for one snapshot. A full overview of the computational time required for each individual step of inference is given in Table IV. Apart from neural network inference, the major time contributions to overall inference time stem from the calculation of fingerprint vectors, LDOS integration, and evaluation of density dependent terms. Parallelization and/or optimization is possible for these calculation steps, with drastic speedups to be expected. Once the cost of creating a ML-DFT model is sufficiently amortized, the computational speedup relative to standard DFT becomes quite stark.

APPENDIX C: SOLID ALUMINUM AT AMBIENT TEMPERATURE AND PRESSURE

We consider modeling of the ambient temperature and pressure aluminum systems as a base case for all future work and experiments. Future investigations of the grid-based ML-DFT model, such as an extensive exploration of alternative ML models, fingerprints, LDOS formulations, and nonaluminum systems, should use this baseline in terms of ML accuracy and computational performance.

In Table V and Fig. 8, we demonstrate that the 298 K model is able to achieve the highest and most consistent accuracy of all models considered in this work. Using a single training snapshot and a single validation snapshot, the ML-DFT model is able to achieve a mean absolute band energy error of 2.48 meV/atom and mean absolute total energy error of 3.07 meV/atom. Furthermore, the 298 K model surpasses the 5 meV/atom accuracy for the most accurate ML-IAPs in large-scale MD simulations [98]. Figure 9 illustrates the predicted DOS in comparison to the reference DOS at 298 K. It confirms our model approximates the DOS with high accuracy. Figure 9 also shows how well the ML model resolves the van Hove singularities in the DOS which are more pronounced at 298 K than at 933 K. We believe this confirms the high accuracy achieved by our ML model. The low variability, seen in Fig. 7, is more easily exploited by ML-DFT's grid-based approach relative to the 933 K cases.

- [1] M. F. Horstemeyer, The near future: ICME for the creation of new materials and structures, in *Integrated Computational Materials Engineering (ICME) for Metals* (Wiley, New York, 2012) Chap. 10, pp. 410–423.
- [2] W. E. Frazier, Metal additive manufacturing: A review, *J. Mater. Eng. Perform.* **23**, 1917 (2014).
- [3] *Advanced Manufacturing and Materials Science: Selected Extended Papers of ICAMMS 2018*, Lecture Notes on Multidisciplinary Industrial Engineering, edited by K. Antony and J. P. Davim (Springer, Berlin, 2018).
- [4] W. L. Mao, H.-k. Mao, P. J. Eng, T. P. Trainor, M. Newville, C.-c. Kao, D. L. Heinz, J. Shu, Y. Meng, and R. J. Hemley, Bonding changes in compressed superhard graphite, *Science* **302**, 425 (2003).
- [5] L. E. Fried, M. R. Manaa, P. F. Pagoria, and R. L. Simpson, Design and synthesis of energetic materials, *Annu. Rev. Mater. Res.* **31**, 291 (2001).
- [6] P. F. Pagoria, G. S. Lee, A. R. Mitchell, and R. D. Schmidt, A review of energetic materials synthesis, *Thermochim. Acta* **384**, 187 (2002).
- [7] M. Werner and W. Fahrner, Review on materials, microsensors, systems and devices for high-temperature and harsh-environment applications, *IEEE Trans. Ind. Electron.* **48**, 249 (2001).
- [8] G. S. Was, *Fundamentals of Radiation Materials Science: Metals and Alloys* (Springer-Verlag, Berlin, 2007).
- [9] R. Graham, B. Morosin, E. Venturini, and M. Carr, Materials modification and synthesis under high pressure shock compression, *Annu. Rev. Mater. Sci.* **16**, 315 (1986).
- [10] A. B. Sawaoka, The role of shock wave compression in materials science, in *Shock Waves in Condensed Matter*, edited by Y. M. Gupta (Springer, Boston, 1986), pp. 51–56.
- [11] D. Alfè and M. J. Gillan, First-Principles Calculation of Transport Coefficients, *Phys. Rev. Lett.* **81**, 5161 (1998).

- [12] B. Militzer, W. B. Hubbard, J. Vorberger, I. Tamblyn, and S. A. Bonev, A massive core in Jupiter predicted from first-principles simulations, *Astrophys. J.* **688**, L45 (2008).
- [13] J. Vorberger, I. Tamblyn, B. Militzer, and S. A. Bonev, Hydrogen-helium mixtures in the interiors of giant planets, *Phys. Rev. B* **75**, 024206 (2007).
- [14] M. Schöttler and R. Redmer, Ab Initio Calculation of the Miscibility Diagram for Hydrogen-Helium Mixtures, *Phys. Rev. Lett.* **120**, 115703 (2018).
- [15] N. Nettelmann, A. Becker, B. Holst, and R. Redmer, Jupiter models with improved *ab initio* hydrogen equation of state (h-reos.2), *Astrophys. J.* **750**, 52 (2012).
- [16] N. Nettelmann, R. Helled, J. Fortney, and R. Redmer, New indication for a dichotomy in the interior structure of Uranus and Neptune from the application of modified shape and rotation data, *Planet. Space Sci.* **77**, 143 (2013).
- [17] W. Lorenzen, B. Holst, and R. Redmer, Demixing of Hydrogen and Helium at Megabar Pressures, *Phys. Rev. Lett.* **102**, 115701 (2009).
- [18] W. Lorenzen, B. Holst, and R. Redmer, Metallization in hydrogen-helium mixtures, *Phys. Rev. B* **84**, 235109 (2011).
- [19] N. Nettelmann, J. J. Fortney, U. Kramm, and R. Redmer, Thermal evolution and structure models of the transiting super-earth GJ 1214b, *Astrophys. J.* **733**, 2 (2011).
- [20] U. Kramm, N. Nettelmann, J. J. Fortney, R. Neuhäuser, and R. Redmer, Constraining the interior of extrasolar giant planets with the tidal Love number k_2 using the example of HAT-P-13b, *Astron. Astrophys.* **538**, A146 (2012).
- [21] B. J. Alder and T. E. Wainwright, Studies in molecular dynamics. I. General method, *J. Chem. Phys.* **31**, 459 (1959).
- [22] S. Plimpton, Fast parallel algorithms for short-range molecular dynamics, *J. Comput. Phys.* **117**, 1 (1995).
- [23] L. A. Zepeda-Ruiz, A. Stukowski, T. Oettel, and V. V. Bulatov, Probing the limits of metal plasticity with molecular dynamics simulations, *Nature (London)* **550**, 492 (2017).
- [24] M. Finnis, *Interatomic Forces in Condensed Matter* (Oxford University Press, Oxford, 2003).
- [25] T. D. Huan, R. Batra, J. Chapman, S. Krishnan, L. Chen, and R. Ramprasad, A universal strategy for the creation of machine learning-based atomistic force fields, *npj Comput. Mater.* **3**, 37 (2017).
- [26] 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 (2017).
- [27] L. Zhang, J. Han, H. Wang, R. Car, and W. E, Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics, *Phys. Rev. Lett.* **120**, 143001 (2018).
- [28] 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**, 136403 (2010).
- [29] A. Jaramillo-Botero, S. Naserifar, and W. A. Goddard, General multiobjective force field optimization framework, with application to reactive force fields for silicon carbide, *J. Chem. Theory Comput.* **10**, 1426 (2014).
- [30] N. Lubbers, J. S. Smith, and K. Barros, Hierarchical modeling of molecular energies using a deep neural network, *J. Chem. Phys.* **148**, 241715 (2018).
- [31] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller, SchNet—a deep learning architecture for molecules and materials, *J. Chem. Phys.* **148**, 241722 (2018).
- [32] A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker, Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials, *J. Comput. Phys.* **285**, 316 (2015).
- [33] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
- [34] N. D. Mermin, Thermal properties of the inhomogeneous electron gas, *Phys. Rev.* **137**, A1441 (1965).
- [35] K. Burke, Perspective on density functional theory, *J. Chem. Phys.* **136**, 150901 (2012).
- [36] A. D. Becke, Perspective: Fifty years of density-functional theory in chemical physics, *J. Chem. Phys.* **140**, 18A301 (2014).
- [37] K. Lejaeghere, G. Bihlmayer, T. Björkman, P. Blaha, S. Blügel, V. Blum, D. Caliste, I. E. Castelli, S. J. Clark, A. D. Corso, S. de Gironcoli, T. Deutsch, J. K. Dewhurst, I. D. Marco, C. Draxl, M. Dułak, O. Eriksson, J. A. Flores-Livas, K. F. Garrity, L. Genovese *et al.*, Reproducibility in density functional theory calculations of solids, *Science* **351**, aad3000 (2016).
- [38] A. Pribram-Jones, D. A. Gross, and K. Burke, DFT: A theory full of holes?, *Annu. Rev. Phys. Chem.* **66**, 283 (2015).
- [39] M. A. Wood, M. A. Cusentino, B. D. Wirth, and A. P. Thompson, Data-driven material models for atomistic simulation, *Phys. Rev. B* **99**, 184305 (2019).
- [40] P. F. Weck, K. R. Cochrane, S. Root, J. M. D. Lane, L. Shulenburg, J. H. Carpenter, T. Sjöström, T. R. Mattsson, and T. J. Vogler, Shock compression of strongly correlated oxides: A liquid-regime equation of state for cerium(IV) oxide, *Phys. Rev. B* **97**, 125106 (2018).
- [41] K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space, *J. Phys. Chem. Lett.* **6**, 2326 (2015).
- [42] 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).
- [43] 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).
- [44] K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K. R. Müller, and E. K. U. Gross, How to represent crystal structures for machine learning: Towards fast prediction of electronic properties, *Phys. Rev. B* **89**, 205118 (2014).
- [45] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, Understanding machine-learned density functionals, *Int. J. Quantum Chem.* **116**, 819 (2016).
- [46] J. C. Snyder, M. Rupp, K. Hansen, L. Blooston, K.-R. Müller, and K. Burke, Orbital-free bond breaking via machine learning, *J. Chem. Phys.* **139**, 224104 (2013).
- [47] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, Finding Density Functionals with Machine Learning, *Phys. Rev. Lett.* **108**, 253002 (2012).

- [48] K. Yao and J. Parkhill, Kinetic energy of hydrocarbons as a function of electron density and convolutional neural networks, *J. Chem. Theory Comput.* **12**, 1139 (2016).
- [49] 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).
- [50] S. R. Broderick and K. Rajan, Eigenvalue decomposition of spectral features in density of states curves, *Europhys. Lett.* **95**, 57005 (2011).
- [51] C. Ben Mahmoud, A. Anelli, G. Csányi, and M. Ceriotti, Learning the electronic density of states in condensed matter, *Phys. Rev. B* **102**, 235130 (2020).
- [52] A. Chandrasekaran, D. Kamal, R. Batra, C. Kim, L. Chen, and R. Ramprasad, Solving the electronic structure problem with machine learning, *npj Comput. Mater.* **5**, 22 (2019).
- [53] A. Abedi, N. T. Maitra, and E. K. U. Gross, Correlated electron-nuclear dynamics: Exact factorization of the molecular wavefunction, *J. Chem. Phys.* **137**, 22A530 (2012).
- [54] M. P. Grumbach, D. Hohl, R. M. Martin, and R. Car, *Ab initio* molecular dynamics with a finite-temperature density functional, *J. Phys.: Condens. Matter* **6**, 1999 (1994).
- [55] A. Alavi, J. Kohanoff, M. Parrinello, and D. Frenkel, *Ab Initio* Molecular Dynamics with Excited Electrons, *Phys. Rev. Lett.* **73**, 2599 (1994).
- [56] A. Alavi, M. Parrinello, and D. Frenkel, *Ab initio* calculation of the sound velocity of dense hydrogen: Implications for models of Jupiter, *Science* **269**, 1252 (1995).
- [57] N. Marzari, D. Vanderbilt, and M. C. Payne, Ensemble Density-Functional Theory for *Ab Initio* Molecular Dynamics of Metals and Finite-Temperature Insulators, *Phys. Rev. Lett.* **79**, 1337 (1997).
- [58] J. L. Alonso, A. Castro, P. Echenique, V. Polo, A. Rubio, and D. Zueco, *Ab Initio* molecular dynamics on the electronic Boltzmann equilibrium distribution, *New J. Phys.* **12**, 083064 (2010).
- [59] A. Abedi, N. T. Maitra, and E. K. U. Gross, Exact Factorization of the Time-Dependent Electron-Nuclear Wave Function, *Phys. Rev. Lett.* **105**, 123002 (2010).
- [60] A. Samanta, M. A. Morales, and E. Schwegler, Exploring the free energy surface using *ab initio* molecular dynamics, *J. Chem. Phys.* **144**, 164101 (2016).
- [61] M. Toda, R. Kubo, and N. Saito, *Statistical Physics I: Equilibrium Statistical Mechanics*, Springer Series in Solid-State Sciences (Springer-Verlag, Berlin, 1983).
- [62] S. Pittalis, C. R. Proetto, A. Floris, A. Sanna, C. Bersier, K. Burke, and E. K. U. Gross, Exact Conditions in Finite-Temperature Density-Functional Theory, *Phys. Rev. Lett.* **107**, 163001 (2011).
- [63] T. Baldsiefen, A. Cangi, and E. K. U. Gross, Reduced-density-matrix-functional theory at finite temperature: Theoretical foundations, *Phys. Rev. A* **92**, 052514 (2015).
- [64] A. Pribram-Jones, S. Pittalis, E. K. U. Gross, and K. Burke, Thermal density functional theory in context, in *Frontiers and Challenges in Warm Dense Matter*, Lecture Notes in Computational Science and Engineering, edited by F. Graziani, M. P. Desjarlais, R. Redmer, and S. B. Trickey (Springer, Cham, 2014), pp. 25–60.
- [65] E. W. Brown, J. L. DuBois, M. Holzmann, and D. M. Ceperley, Exchange-correlation energy for the three-dimensional homogeneous electron gas at arbitrary temperature, *Phys. Rev. B* **88**, 081102(R) (2013).
- [66] V. V. Karasiev, T. Sjöström, J. Dufty, and S. B. Trickey, Accurate Homogeneous Electron Gas Exchange-Correlation Free Energy for Local Spin-Density Calculations, *Phys. Rev. Lett.* **112**, 076403 (2014).
- [67] S. Groth, T. Dornheim, T. Sjöström, F. D. Malone, W. M. C. Foulkes, and M. Bonitz, *Ab Initio* Exchange-Correlation Free Energy of the Uniform Electron Gas at Warm Dense Matter Conditions, *Phys. Rev. Lett.* **119**, 135001 (2017).
- [68] P. E. Blöchl and M. Parrinello, Adiabaticity in first-principles molecular dynamics, *Phys. Rev. B* **45**, 9413 (1992).
- [69] G. Galli, R. M. Martin, R. Car, and M. Parrinello, *Ab initio* calculation of properties of carbon in the amorphous and liquid states, *Phys. Rev. B* **42**, 7470 (1990).
- [70] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [71] R. M. Dreizler and E. K. U. Gross, *Density Functional Theory* (Springer, Berlin, 1990).
- [72] J. Behler and M. Parrinello, Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [73] A. Seko, A. Takahashi, and I. Tanaka, Sparse representation for a potential energy surface, *Phys. Rev. B* **90**, 024101 (2014).
- [74] J. C. Snyder, M. Rupp, K.-R. Müller, and K. Burke, Nonlinear gradient denoising: Finding accurate extrema from inaccurate functional derivatives, *Int. J. Quantum Chem.* **115**, 1102 (2015).
- [75] 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).
- [76] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. D. Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos *et al.*, QUANTUM ESPRESSO: A modular and open-source software project for quantum simulations of materials, *J. Phys.: Condens. Matter* **21**, 395502 (2009).
- [77] P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. B. Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. D. Corso, S. de Gironcoli, P. Delugas, R. A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo *et al.*, Advanced capabilities for materials modelling with Quantum ESPRESSO, *J. Phys.: Condens. Matter* **29**, 465901 (2017).
- [78] P. Giannozzi, O. Baseggio, P. Bonfà, D. Brunato, R. Car, I. Carnimeo, C. Cavazzoni, S. de Gironcoli, P. Delugas, F. Ferrari Ruffino, A. Ferretti, N. Marzari, I. Timrov, A. Urru, and S. Baroni, Quantum ESPRESSO toward the exascale, *J. Chem. Phys.* **152**, 154105 (2020).
- [79] D. R. Hamann, Optimized norm-conserving Vanderbilt pseudopotentials, *Phys. Rev. B* **88**, 085117 (2013).
- [80] J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou, and K. Burke, Restoring the Density-Gradient Expansion for Exchange in Solids and Surfaces, *Phys. Rev. Lett.* **100**, 136406 (2008).
- [81] H. J. Monkhorst and J. D. Pack, Special points for Brillouin-zone integrations, *Phys. Rev. B* **13**, 5188 (1976).
- [82] L. Van Hove, The occurrence of singularities in the elastic frequency distribution of a crystal, *Phys. Rev.* **89**, 1189 (1953).
- [83] P. E. Blöchl, O. Jepsen, and O. K. Andersen, Improved tetrahedron method for Brillouin-zone integrations, *Phys. Rev. B* **49**, 16223 (1994).

- [84] R. Drautz, Atomic cluster expansion for accurate and transferable interatomic potentials, *Phys. Rev. B* **99**, 014104 (2019).
- [85] J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, Recent advances and applications of machine learning in solid-state materials science, *npj Comput. Mater.* **5**, 83 (2019).
- [86] A. Mathuriya, D. Bard, P. Mendygral, L. Meadows, J. Arnemann, L. Shao, S. He, T. Kärmä, D. Moise, S. J. Pennycook, K. Maschhoff, J. Sewall, N. Kumar, S. Ho, M. F. Ringenburt, Prabhat, and V. Lee, Cosmoflow: Using deep learning to learn the universe at scale, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC '18* (IEEE, New York, 2018).
- [87] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2016), pp. 770–778.
- [88] A. Ziletti, D. Kumar, M. Scheffler, and L. Ghiringhelli, Insightful classification of crystal structures using deep learning, *Nat. Commun.* **9**, 2775 (2018).
- [89] T. Young, D. Hazarika, S. Poria, and E. Cambria, Recent trends in deep learning based natural language processing [review article], *IEEE Comput. Intell. Mag.* **13**, 55 (2018).
- [90] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (Association for Computational Linguistics, Minneapolis, MN, 2019), pp. 4171–4186.
- [91] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, *Neural Networks* **6**, 861 (1993).
- [92] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016), <http://www.deeplearningbook.org>.
- [93] L. Prechelt, Early stopping—but when?, in *Neural Networks: Tricks of the Trade* (Springer, 1996), pp. 55–69.
- [94] R. M. Lewis, V. Torczon, and M. W. Trosset, Direct search methods: Then and now, *J. Comput. Appl. Math.* **124**, 191 (2000), Numerical Analysis 2000, Vol. IV: Optimization and Nonlinear Equations.
- [95] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, Automatic differentiation in PyTorch, Neural Information Processing System (2017).
- [96] T. E. Oliphant, *A Guide to NumPy*, Vol. 1 (Trelgol Publishing, 2006), 2nd ed.
- [97] A. Sergeev and M. D. Balso, Horovod: Fast and easy distributed deep learning in TensorFlow, [arXiv:1802.05799](https://arxiv.org/abs/1802.05799).
- [98] Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csanyi, A. V. Shapeev, A. P. Thompson, M. A. Wood, and S. P. Ong, Performance and cost assessment of machine learning interatomic potentials, *J. Phys. Chem. A* **124**, 731 (2020).
- [99] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in PYTHON, *J. Mach. Learn. Res.* **12**, 2825 (2011).