

ARTICLE OPEN



Linear Jacobi-Legendre expansion of the charge density for machine learning-accelerated electronic structure calculations

Bruno Focassio^{1,2,3}, Michelangelo Domina³, Urvesh Patil³, Adalberto Fazio^{1,2} and Stefano Sanvito³

Kohn–Sham density functional theory (KS-DFT) is a powerful method to obtain key materials' properties, but the iterative solution of the KS equations is a numerically intensive task, which limits its application to complex systems. To address this issue, machine learning (ML) models can be used as surrogates to find the ground-state charge density and reduce the computational overheads. We develop a grid-centred structural representation, based on Jacobi and Legendre polynomials combined with a linear regression, to accurately learn the converged DFT charge density. This integrates into a ML pipeline that can return any density-dependent observable, including energy and forces, at the quality of a converged DFT calculation, but at a fraction of the computational cost. Fast scanning of energy landscapes and producing starting densities for the DFT self-consistent cycle are among the applications of our scheme.

npj Computational Materials (2023)9:87 ; <https://doi.org/10.1038/s41524-023-01053-0>

INTRODUCTION

As the main workhorse in electronic structure calculations, density functional theory (DFT)^{1,2} is today the most widely used method to compute materials properties. Its success derives from the favourable trade-off between computational overheads and accuracy, even when using simple approximations for the exchange and correlation energy functional^{2–5}. The central quantity in DFT is the electron charge density that, in principle, gives access to the ground-state properties¹, and of particular interest, to the ground-state total energy. In practice, however, the DFT functional is never minimized directly by using the charge density⁶, but rather by solving a self-consistent set of single-particle equations, known as the Kohn–Sham (KS) equations². This procedure effectively imposes a computational bottleneck and although large-scale calculations can be performed^{7,8}, the typical system routinely simulated by DFT rarely reaches a few hundred atoms.

Machine learning (ML) has recently emerged as a surrogate for solving DFT KS equations and possibly replacing them^{9–11}. For instance, trained ML models can be used as predictors for properties such as the energy gaps^{12–14}, superconducting critical temperatures^{15–19}, thermodynamic stability²⁰, topological invariant^{21,22}, just to name a few. These models learn a direct map between the structure/composition and the target property, thus avoiding one or many computationally expensive calculations. Using ML for such mapping comes at the cost of accuracy, transferability, physical insight and the need for a large volume of high-quality training data, usually obtained through these very same computationally expensive calculations or, more rarely, from experimental sources^{23,24}.

For tasks such as structure prediction^{25–28}, phase diagrams evaluation^{29–31}, molecular dynamics^{32–34}, and, more generally, materials discovery^{35–38} one requires fast access to accurate energy, forces and stress tensor of the system investigated. Machine learning inter-atomic potentials (ML-IAPs) are developed to this end, bridging the gap between ab initio methods and empirical force fields. The several strategies proposed to date

implement a diversity of structural representations and learning algorithms^{39–44} to design ML-IAPs attaining accuracies close to that of DFT at a small fraction of the computational cost⁴³. The performance of these models is not only a product of the representation of the atomic structure and the ML algorithm, but also the volume, quality, and diversity of the data play a fundamental role^{43,45}. In general, the construction of ML-IAPs requires campaigns of DFT calculations, whose extension and quality depend on the problem at hand (e.g., the number of species present in a given compound) and the range of applicability of the potential (e.g., the temperature range).

A radically different use of ML consists of improving the theory at its core instead of targeting the DFT outputs. For instance, ML can be used to numerically design new energy density functionals, effectively producing fully exchange and correlation energies^{46–51}, and kinetic energy densities^{52–55}. These strategies, in general, seek to find more accurate approximations to the DFT energy, going beyond the current approximations⁵⁶, or to eliminate the need of introducing the KS construct by replacing the self-consistent KS equations with a direct minimization of the functional⁶. Unfortunately, although promising, these approaches are still far from obtaining a “universal” functional, treating all systems on an equal footing⁵⁶. Note also that the construction of ML functionals requires results obtained at the wave-function quantum-chemistry level, a highly computationally expensive task.

In the same spirit, an alternative way to include ML in the DFT workflow is to construct models to directly predict the converged target DFT quantities, namely the Hamiltonian⁵⁷, the wavefunctions^{58,59}, and the electron density^{60–65}. The goal here is not that of improving the functional, but to reduce or completely eliminate the number of iterative steps needed to solve the KS equations. There are two main approaches used to predict the electronic charge density, $n(\mathbf{r})$, through ML. One possibility is to expand $n(\mathbf{r})$ over a local-orbital basis set and learn by ML the expansion coefficients. The completeness of such expansion, the basis set details, and the size of the training data limit the accuracy of the ML model^{60,61} and may introduce errors intrinsic to the particular

¹Federal University of ABC (UFABC), Santo André, São Paulo 09210-580, Brazil. ²Ilum School of Science, CNPEM, Campinas, São Paulo 13083-970, Brazil. ³School of Physics and CRANN Institute, Trinity College Dublin, Dublin 2, Ireland. ✉email: b.focassio@ufabc.edu.br; sanvitos@tcd.ie

representation⁶². Also, the approach is not transferable, namely a different ML model must be constructed for any different basis set.

The second approach considers the real-space representation of $n(\mathbf{r})$, which is written over a grid in Cartesian space. This is a more “natural” representation available in any DFT code. Its main advantage is that the value of the electron density at a grid point is rotationally invariant with respect to the external potential, namely with respect to the position of the surrounding nuclei. As such, one can construct ML models that predict $n(\mathbf{r})$ one grid point at the time, using as descriptors the local atomic neighbourhood of any given grid point (within some chosen cutoff radius). The success of such a grid-based approach largely depends on the chosen representation for the local environment and the learning algorithm. Usually, a single DFT calculation results in tens of millions of grid points so that the generation of abundant training data appears like an easy computational task. However, in a single calculation, there is data redundancy and little diversity (a narrow distribution of external potentials is explored), so multiple configurations for the same systems are usually considered. Then, one typically constructs large neural networks with millions of weights to be learned^{63,64}, resulting in generally heavy models with little transferability.

Here our main focus is to transform such a grid-based approach into a lightweight tool that can be universally applied to DFT calculations. This is achieved by drastically reducing the computational overheads while reaching extremely high accuracies. In particular, we introduce a grid-centred representation of the atomic structure based on the Jacobi and Legendre (JL) polynomials, which were previously proposed to construct efficient ML force fields⁶⁶. The JL representation is used to build a linear regression for the charge density, where the many-body contributions of different orders are separated. This results in a very compact model with a few coefficients to be trained on a small subset of the total number of grid points available. For the sake of brevity, we call such a class of models Jacobi-Legendre charge density models (JLCDMs). The efficiency and accuracy of our scheme are demonstrated for a range of molecules and solids, including benzene, aluminium, molybdenum, and two-dimensional MoS₂. In particular, we show that the KS self-consistent cycle can be bypassed completely in calculating fully converged total energies and forces. Our method is implemented to work with the widely used Vienna ab initio simulation package (VASP)^{67,68}.

RESULTS

Figure 1 provides a schematic view of the construction of a JLCDM. Given an atomic configuration, the space is subdivided

into a Cartesian grid, and the atomic environment (the position of the atoms) of each grid is described by an expansion of JL polynomials. A selected number of such expansions forms the training set of a linear regression model that predicts the charge density over the entire grid. Finally, this is used as the converged ground-state density to evaluate energy, forces, and any other density-dependent observable, $\langle O \rangle$, or as a starting point for self-consistent KS-DFT calculations.

Linear expansion of the charge density

The charge density, $n(\mathbf{r})$, at a grid point \mathbf{r}_g can be separated into many-body contributions as

$$n(\mathbf{r}_g) = n^{(1)}(\mathbf{r}_g) + n^{(2)}(\mathbf{r}_g) + n^{(3)}(\mathbf{r}_g) + \dots + n^{(m)}(\mathbf{r}_g) \quad (1)$$

where $n^{(m)}$ is the m^{th} -body (mB) term of the expansion. Thus, $n^{(1)}(\mathbf{r}_g)$ encodes the atomic contributions to the charge density at \mathbf{r}_g , $n^{(2)}$ is the contribution from atom pairs, $n^{(3)}$ is the contribution from atoms triplets, etc. Equation (1) can then be rewritten as

$$n(\mathbf{r}_g) = \sum_i n_i^{(1)}(\mathbf{r}_g) + \sum_{i \neq j} n_{ij}^{(2)}(\mathbf{r}_g) + \sum_{i \neq j, i \neq k, j \neq k} n_{ijk}^{(3)}(\mathbf{r}_g) + \dots \quad (2)$$

where the sums over the $ij,k\dots$ indexes run over the atoms neighbouring the grid point at \mathbf{r}_g up to the cutoff distance, r_{cut} . The assumption that the electron density at one point is determined mostly by the external potential generated by the closest atoms follows from the wave mechanics’ locality principle⁶⁹.

The atomic configurations required by each contribution in the expansion are expressed through a local representation that here we generally call “fingerprint”. The fingerprints should be: (i) invariant by translations, (ii) invariant by global rotations of the atoms in the reference frame of the grid point, (iii) invariant to changes in the coordinate system, (iv) invariant to permutations of the atomic indices. Furthermore, they should provide a continuous map of the atomic neighbourhood, i.e., small changes in the atomic structure must reflect small changes in the fingerprints. Finally, the fingerprints should be uniquely determined⁷⁰ and computationally cheap.

Following closely ref. ⁶⁶, we expand the one-body contribution, $n_i^{(1)}$, using the distances between the grid point and the atomic

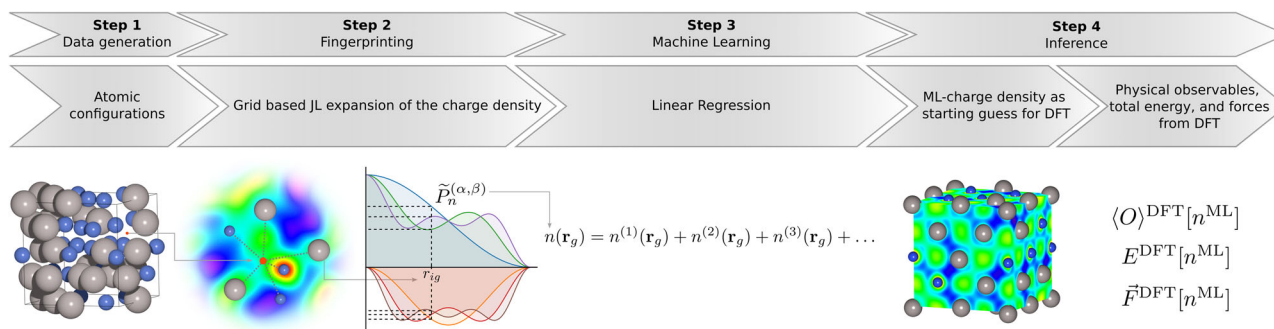


Fig. 1 Illustration of the workflow used to construct a JLCDM predicting the converged DFT ground-state charge density and the associated observables. (Step 1) The procedure starts with an atomic distribution and the mapping of the space over a Cartesian grid. (Step 2) Each grid point is associated with a local atomic environment described by the Jacobi-Legendre expansion. Such expansion is used to construct a linear model (Step 3) that, once trained, accurately predicts the charge density of the grid point. After computing the charge density over the entire grid, this is used to perform DFT calculations (Step 4). For instance, the total energy, the atomic forces, and other density-dependent observables can be easily obtained by using a few steps of frozen-density KS-DFT instead of the full self-consistent cycle.

neighbourhood as

$$n_i^{(1)}(\mathbf{r}_g) = \sum_{n=1}^{n_{\max}} a_n^{Z_i} \tilde{P}_n^{(\alpha,\beta)} \left(\cos \left(\pi \frac{r_{ig} - r_{\min}}{r_{\text{cut}} - r_{\min}} \right) \right) \quad (3)$$

$$\tilde{P}_n^{(\alpha,\beta)}(x) = \begin{cases} P_n^{(\alpha,\beta)}(x) - P_n^{(\alpha,\beta)}(-1) & \text{for } -1 \leq x \leq 1 \\ 0 & \text{for } x < -1 \end{cases} \quad (4)$$

with $P_n^{(\alpha,\beta)}$ being the Jacobi polynomial of order n . Here, $r_{ig} = |\mathbf{r}_i - \mathbf{r}_g|$ is the distance between the grid point g at \mathbf{r}_g and the i th atom i at \mathbf{r}_i , r_{cut} is the radius cutoff, r_{\min} is a distance shift parameter in the range $(-\infty, r_{\text{cut}})$. The degree of the expansion is set by n with the sum running in the interval $[1, n_{\max}]$, while α and β control the shape of the polynomial with $\alpha, \beta > -1$. Note that the choice of the basis used to expand the atomic structure is not unique. Jacobi polynomials represent a convenient one, since they effectively describe a vast class of basis-set types. For instance $(\alpha, \beta) = (0, 0)$ gives us the Legendre polynomials, while $(\alpha, \beta) = (0.5, 0.5)$ the Chebyshev polynomials of the second kind. Here we decide to maintain the generality and treat the (α, β) pair as an hyperparameter to optimize. The expansion coefficients $a_n^{Z_i}$ in Eq. (3) depend on the atomic species considered. As defined in Eq. (4), the “vanishing Jacobi polynomials” smoothly vanish at the cutoff radius without needing an additional ad hoc cutoff function.

The terms forming the two-body contribution, $n_{ij}^{(2)}$, can be uniquely written as a function of two distances, r_{ig} and r_{jg} , and the cosine of the subtended angle at g , $\hat{\mathbf{r}}_{ig} \cdot \hat{\mathbf{r}}_{jg}$. We then expand the distances over the vanishing Jacobi polynomials and the angle over Legendre polynomials. The expansion can then be written as,

$$n_{ij}^{(2)}(\mathbf{r}_g) = \sum_{n_1, n_2=1}^{n_{\max}} \sum_{l=0}^{l_{\max}} a_{n_1 n_2 l}^{Z_i Z_j} \tilde{P}_{n_1 i g}^{(\alpha,\beta)} \tilde{P}_{n_2 j g}^{(\alpha,\beta)} P_l^{jig}, \quad (5)$$

where we have used the shorthand notations

$$\tilde{P}_{nig}^{(\alpha,\beta)} = \tilde{P}_n^{(\alpha,\beta)} \left(\cos \left(\pi \frac{r_{ig} - r_{\min}}{r_{\text{cut}} - r_{\min}} \right) \right),$$

and $P_l^{jig} = P_l(\hat{\mathbf{r}}_{ig} \cdot \hat{\mathbf{r}}_{jg})$, P_l is the Legendre polynomial, $\hat{\mathbf{r}}_{pg} = (\mathbf{r}_p - \mathbf{r}_g)/r_{pg}$, and l defines the Legendre expansion degree with the sum running in the interval $[0, l_{\max}]$. The Legendre polynomials are the natural choice for expanding the scalar products between two real space vectors in three dimensions, as suggested by the addition theorem of spherical harmonics. As in the one-body case, the expansion coefficients $a_{n_1 n_2 l}^{Z_i Z_j}$ depend on the pair of atomic species considered. The Jacobi indices n_1 and n_2 , and the atom indices i and j are symmetric under the simultaneous swap, therefore if $Z_i = Z_j$ only terms $n_1 \geq n_2$ should be considered.

Notice that, in the m -body expansion for $m > 1$, angular information enters via a pairwise dot product of unit vectors joining the atoms to the grid point. The unit vectors are ill-defined when the distance of the grid point from the atom approaches zero and creates a discontinuity in the fingerprints. Assuming that the atomic contribution (1B term) to the charge density dominates at very small distances from the nucleus, we can introduce a double-vanishing Jacobi polynomial in place of the simple vanishing one for all the m -body expansions with $m > 1$ as given in Eqs. (7) and (8). The double-vanishing Jacobi polynomials are defined as

$$\bar{P}_n^{(\alpha,\beta)}(x) = \tilde{P}_n^{(\alpha,\beta)}(x) - \frac{\tilde{P}_n^{(\alpha,\beta)}(1)}{\tilde{P}_1^{(\alpha,\beta)}(1)} \tilde{P}_1^{(\alpha,\beta)}(x) \text{ for } n \geq 2 \quad (6)$$

with $x = \cos \left(\pi \frac{r_{ig} - r_{\min}}{r_{\text{cut}} - r_{\min}} \right)$. Equation (5) now reads

$$n_{ij}^{(2)}(\mathbf{r}_g) = \sum_{n_1, n_2=2}^{n_{\max}} \sum_{l=0}^{l_{\max}} a_{n_1 n_2 l}^{Z_i Z_j} \bar{P}_{n_1 i g}^{(\alpha,\beta)} \bar{P}_{n_2 j g}^{(\alpha,\beta)} P_l^{jig} \quad (7)$$

with $n_1, n_2 \geq 2$. Generally, a m -body cluster centred on the grid point g can be uniquely defined by m distances and the $m(m-1)/2$

angles subtended at g . Using the recipe from Eqs. (3) and (7), the m -body expansion can then be written by associating a Jacobi polynomial to each distance and a Legendre polynomial to each angle. For instance, the three-body contribution $n_{ijk}^{(3)}$ is of the form

$$n_{ijk}^{(3)}(\mathbf{r}_g) = \sum_{n_1, n_2, n_3=2}^{n_{\max}} \sum_{l_1, l_2, l_3}^{l_{\max}} a_{n_1 n_2 n_3 l_1 l_2 l_3}^{Z_i Z_j Z_k} \times \bar{P}_{n_1 i g}^{(\alpha,\beta)} \bar{P}_{n_2 j g}^{(\alpha,\beta)} \bar{P}_{n_3 k g}^{(\alpha,\beta)} P_{l_1}^{jig} P_{l_2}^{jkg} P_{l_3}^{ikg} \quad (8)$$

Using this charge density expansion at each grid point, we can generate a linear representation of the charge density in the expansion coefficients. Therefore, we can learn the ground state charge density by using linear regression, as

$$n^{\text{DFT}}(\mathbf{r}_g) = \sum_i \sum_n a_n^{Z_i} \tilde{P}_{nig}^{(\alpha,\beta)} + \sum_{i \neq j} \sum_{n_1 n_2} \sum_{l=0}^{l_{\max}} a_{n_1 n_2 l}^{Z_i Z_j} \bar{P}_{n_1 i g}^{(\alpha,\beta)} \bar{P}_{n_2 j g}^{(\alpha,\beta)} P_l^{jig} + \sum_{\substack{i \neq j \\ i \neq k \\ j \neq k}} \sum_{n_1 n_2 n_3} \sum_{l_1 l_2 l_3}^{l_{\max}} a_{n_1 n_2 n_3 l_1 l_2 l_3}^{Z_i Z_j Z_k} \bar{P}_{n_1 i g}^{(\alpha,\beta)} \bar{P}_{n_2 j g}^{(\alpha,\beta)} \bar{P}_{n_3 k g}^{(\alpha,\beta)} P_{l_1}^{jig} P_{l_2}^{jkg} P_{l_3}^{ikg} + \dots \quad (9)$$

In the next section, we will demonstrate the prediction power of our JLCDM for a benzene molecule, for periodic solids such as aluminium (Al) and molybdenum (Mo), and a two-dimensional material MoS₂. We will also demonstrate the generalization power of JLCDM for previously unknown phases of Al and MoS₂. Finally, we will show that the charge density predicted by our model can be fed back into popular DFT codes to accurately calculate the total energy and forces at a fraction of the typical numerical cost.

Grid-point sampling strategy

We start our analysis by discussing the construction of an appropriate training set for our JLCDM, which is truncated at the 2-body order since this is already enough for extremely accurate predictions. Previously published works^{63,64} have trained large neural networks over the entire grid-point mesh, typically containing a few million density values. Here we show that this is not necessary since there is significant redundancy in the information, and often the inclusion of the entire density in the training set has just the effect of producing an unbalanced ensemble. This is easy to see in the case of molecules, where most grid points are situated far away from the molecule and, by sitting in a vacuum, possess similar vanishing small charge density. For this reason, we implement a sampling strategy that allows us to use only a small fraction of the grid points but includes more diverse atomic arrangements.

In practice, our simple sampling scheme consists in assigning to a point \mathbf{r} in space a probability of selection based on the value of the charge density, $n(\mathbf{r})$, at that point. The probability of selection is given by a normal distribution of the inverse of the charge density, namely $\exp[-(1/n(\mathbf{r}))^2/2\sigma^2]$. This choice gives more importance to grid points presenting large electron densities, while low-density regions will contribute little to the training set. The parameter σ controls how sharp or broad this probability distribution is, a tool that helps us to select grid points closer or farther away from the charge density maxima. Such a targeted sampling technique is accompanied by uniform sampling across the unit cell, which guarantees that enough diversity is maintained in the training set. As a result, we can construct an accurate model trained with just about 0.1% of the available training points (see the “Methods” section for more details). Note

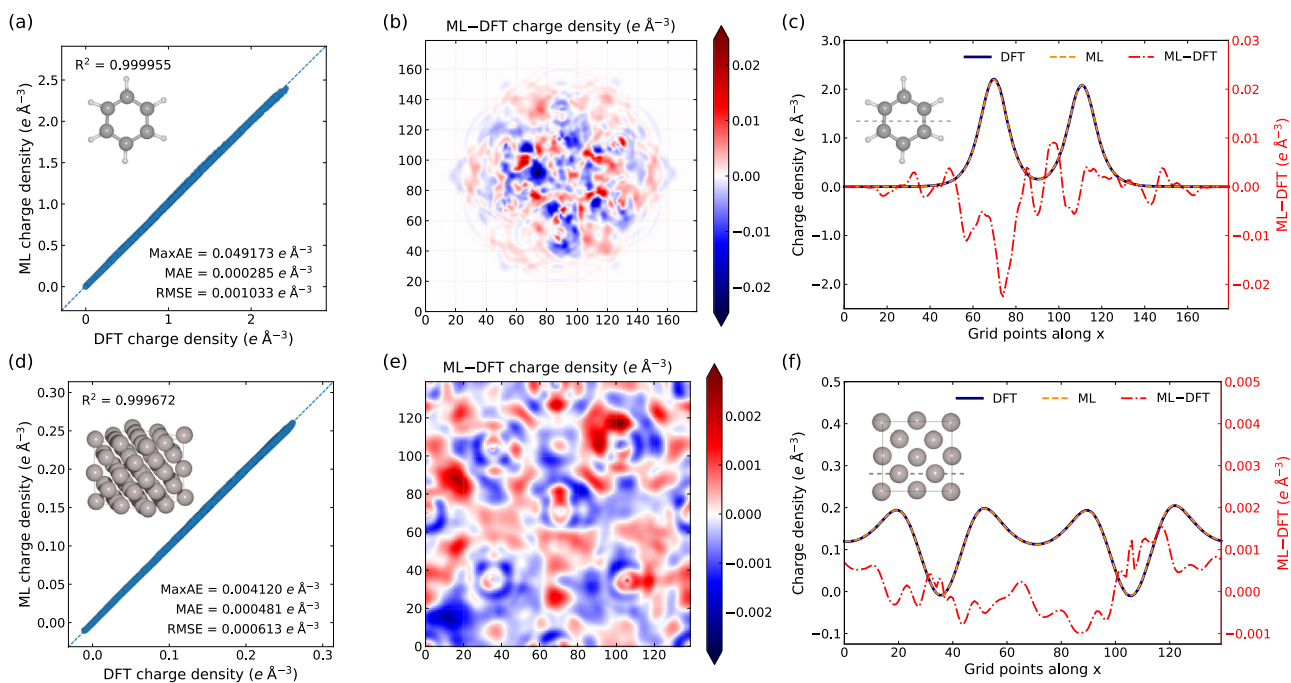


Fig. 2 Analysis of the performance of the JLCDM. **a** Displays the parity plot for the benzene test set together with mean absolute error (MAE), root-mean-squared error (RMSE), maximum absolute error (MaxAE) and R^2 metrics. **b** Displays the charge density difference (the error) between the fully converged DFT ground-state density and that predicted by the model for a distorted benzene configuration selected from the test set. The results for a symmetric benzene molecule can be found in the Supplemental Information (SI). Here we show the plane containing the molecule. **c** Shows DFT and JLCDM-predicted charge density for benzene computed along the line indicated in the inset. The plot also reports their difference with values provided on the right-hand side scale (red). **d** Displays the parity plot for the aluminium test set. **e** Displays the charge density difference (error) between the fully converged DFT ground-state density and that predicted by the model for a distorted aluminium configuration selected from the test set. The slice shows the basal plane of the supercell ($z = 0.0 \text{ \AA}$). In **f**, the DFT and ML charge density for aluminium is computed along the line indicated in the inset. The plot also reports their difference with values provided on the right-hand side scale (red). The planes chosen in **c**, **f** are the same as those in **b**, **e**, respectively.

that our sampling strategy is not limited to linear charge density expansions. The same can be used as an efficient way to train even neural network models, resulting in much smaller models attaining the same or higher accuracy.

Accuracy of the models

We now discuss the accuracy that can be reached by the JLCDMs for both molecules and solids. Figure 2a displays the parity plot of the charge density at the grid points for the 30 atomic configurations contained in the test set of the benzene molecule. These have been obtained from ref. ⁶² by molecular dynamics at 300 K and performing DFT calculations on each sampled geometry. For benzene, our 1B + 2B JLCDM contains 1572 coefficients trained over 6,000 density-grid points, out of the 5,832,000 available per atomic configuration over the 30 configurations used for training and another 30 for testing. The test-set mean absolute error (MAE) achieved is $0.000285 \text{ e \AA}^{-3}$. Such error corresponds to 0.011% of the maximum density, meaning that the charge density obtained by the JLCDM is very close to that of a well-converged DFT calculation. Note that the MAE on the total electron count is 0.025. The model and sampling hyperparameters are reported in Tables 3 and 5, respectively, in the “Methods” section.

In Fig. 2b, we present the difference between the charge density obtained with JLCDM and the converged DFT charge density on a plane, while Fig. 2c shows a line scan in the same plane of the two charge densities and their difference. As expected, the absolute error is larger in the region closer to the nuclei, where the charge density is maximized. However, no emerging pattern indicates that the JLCDM is not biased against any particular local atomic configuration. Importantly, the error, as

the density, vanishes for positions far from the molecule. Our constructed JLCDM performs better than published models⁶³ despite being trained over a tiny fraction of data and being constructed on only 1572 trainable parameters.

Next, we move to metallic solids, aluminium and molybdenum. Aluminium is a benchmark system chosen for comparison with previously published models^{61,63,64}. Its electronic structure features a very delocalized charge density, so as a second example, we also consider an early transition metal, Mo, which presents a higher degree of charge localization. In constructing the JLCDM, we use the same density sampling procedure as used for the benzene molecule. See the “Methods” section for details.

In the case of Al, we train and test over 10 configurations obtained from ab initio molecular dynamics (AIMD). We find that a 1B + 2B JLCDM with only 120 trainable coefficients gives us a MAE of $0.000481 \text{ e \AA}^{-3}$, at par with previously published deep neural networks^{63,64}. Most importantly, our model generalizes better, as we will show in the next section. Figure 2d shows the parity plot for the Al test set, demonstrating the accuracy achieved. Similarly to the case of benzene, the difference between the ML predictions and the DFT charge density does not present any clear error pattern, see Fig. 2e, except for the expected increase close to the nuclei. In general, the charge density error for Al is found to be 10 times smaller than that found for benzene, as one can see from the line plot of Fig. 2f.

Similar results are also obtained for Mo, where a JLCDM with 812 trainable parameters returns a MAE and a RMSE of $0.001974 \text{ e \AA}^{-3}$ and $0.002820 \text{ e \AA}^{-3}$, respectively, see the Supplemental Information (SI) for details. In contrast to benzene and aluminium, the charge density error appears to have a radial distribution centred around each atom with a minimum error in the interstitial region. The maximum absolute error over the test




Fig. 3 Analysis of the performance of the JLCDM for MoS₂. **a** Shows the parity plot between JLCDM-predicted and DFT charge density for the three MoS₂ polymorphs, 1H, 1T, and 1T'. In this case 1H and 1T phases are used for training, while the model is tested on the 1T'. All the error metrics shown, R², MAE, RMSE, and MaxAE, correspond to the test set. The inset depicts a snapshot of 1T'-MoS₂. **b** Display the charge density difference (error) between the fully converged DFT charge density and the JLCDM-predicted one over the plane of the monolayer ($z = c/2$ Å) for a distorted 1T'-MoS₂ configuration selected from the test set. **c** Shows the charge density profile for fully converged DFT and JLCDM-predicted charge density along the path highlighted with a dashed line in **b**. The difference between charge densities can be read on the right-hand side scale (red) of **c**.

set in this case is only ~ 0.06 eÅ⁻³, and it is found over a small set of grid points.

Finally, we focus on two-dimensional MoS₂, which helps us to demonstrate the capability of our JLCDM to generalize to previously unseen phases. Two-dimensional MoS₂ can be found in multiple polymorphs, both semiconducting and metallic. Also, for MoS₂, we use the same charge-density sampling procedure adopted for benzene, Al and Mo. See the “Methods” section. However, this time we train and test the model on different phases; namely, the training set is constructed using atomic configurations of the 1H and 1T phases while we test our prediction on the 1T' phase. The 1H phase is formed by sandwiched hexagonal layers of S–Mo–S in a Bernal stacking, while the 1T phase presents a rhombohedral arrangement⁷¹. As the free-standing 1T phase is unstable, a spontaneous lattice distortion in the x direction creates the 1T' one^{71,72}, which is depicted in the inset of Fig. 3a. The three polymorphs present completely different electronic structures. The 1H phase is semiconducting with a 1.58 eV theoretical energy gap, while the 1T phase is metallic⁷³. In contrast, the 1T' polymorph has a topological gap (0.08 eV) induced by spin-orbit coupling⁷⁴, while it remains a semi-metal in the absence of spin-orbit coupling interaction.

In order to train our JLCDM, we use 10 AIMD (at 300 K) configurations each for the 1H and 1T phases, while the test set is made of ten 1T' AIMD (at 300 K) snapshots. Figure 3a shows the parity plot for all three polymorphs, namely for the training and test set. By visual inspection, one can notice that the error slightly increases for the 1T' phase, but the JLCDM still performs extremely well, displaying a MAE and a RMSE of 0.002725 eÅ⁻³ and 0.008080 eÅ⁻³, respectively. Also, the JCDM remains compact with 2,346 trainable parameters in this case. The charge density difference plot, see Fig. 3b, tells us that the error tends to be larger in the region around the Mo ions pointing towards the S atoms. This feature is somehow expected since the bonding structure of the three phases is different, trigonal prismatic for 1H, octahedral for 1T phase, and a distorted lattice for 1T'. The line density plot of Fig. 3c further shows that the JLCDM slightly overestimates the charge density surrounding the Mo atoms. However, it is worth noting that the error is small, <2%, so the JLCDM-predicted charge density for the unseen 1T' phase is still of high quality, namely, the JLCDM can be used to explore new phases.

JLCDM performance on the DFT total energy and forces

In the previous section, we have shown that the charge density predicted by our JLCDM is close to the DFT converged one. Now

Table 1. JLCDM performance metrics on the task of predicting total energy and forces.

System	Total energy (meV per atom)		Forces (eV Å ⁻¹)	
	MAE	RMSE	MAE	RMSE
Benzene	4.021	4.065	0.031	0.046
Al	0.046	0.054	0.007	0.009
Mo	0.203	0.212	0.019	0.024
1T'-MoS ₂	8.058	8.845	0.078	0.104

These are obtained through non-self-consistent DFT using the JLCDM-predicted charge density. The force error is computed over all and all atoms.

we show that the energy and forces corresponding to such charge density are close to the corresponding converged values, with the average error matching those of state-of-the-art machine-learning force fields.

This is demonstrated by constructing the KS Hamiltonian corresponding to the JLCDM-predicted charge density. The band energy contribution to the total energy, $E_{\text{band}} = \sum_i f(\epsilon_i) \epsilon_{i\mu}$ is obtained by summing up the occupied KS eigenvalues, ϵ_i [$f(\epsilon_i)$ is the occupation number], which are computed by diagonalizing the KS Hamiltonian. The remaining contributions to the total energy are obtained directly from the JLCDM electron density. Such a scheme is implemented in the VASP code, where an interactive matrix-diagonalization procedure requires performing a set of non-self-consistent iterations to compute the KS eigenvalues and eigenvectors, i.e. the charge density is not updated during these iterations. In this work, we select 5 non-self consistent iterations for all systems. The total energy and forces obtained are then compared with those computed through a converged fully self-consistent DFT calculation. As given by such procedure, the total energy yielded by the JLCDM-predicted charge density may be lower than the KS-DFT ground-state energy.

The MAE and RMSE metrics of the calculated energy and forces are given in Table 1, while Fig. 4 shows the error distributions as box and violin plots. Aluminium presents the narrower total-energy error spread, with values ranging from -0.11 meV per atom to -0.02 meV per atom and with a mean error at -0.05 meV per atom. This is then followed by Mo, with a total-energy error spread between 0.12 meV per atom and 0.33 meV per atom with a mean error at 0.20 meV per atom, and then benzene, with a total-

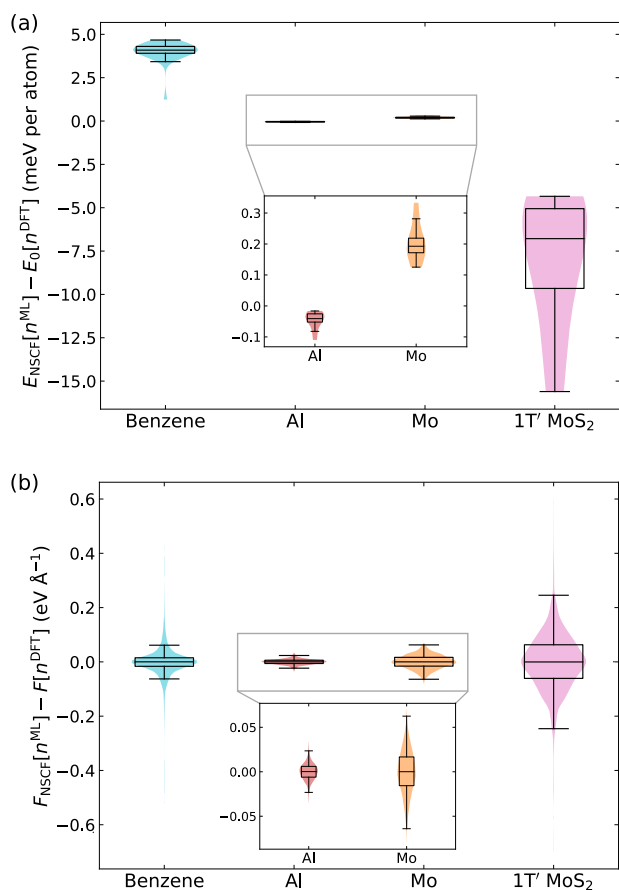


Fig. 4 Error on the total energy and forces. Box and violin plots for the error on the total energy (a) and the forces (b) computed from JLCDM-predicted charge density. The fully converged DFT values provide the ground truth. The insets show a magnified version of the results for Al and Mo, whose distributions are very narrow on the global scale. The associated absolute mean values are reported in Table 1. The lines in the middle of the boxes mark the medians. The boxes are plotted from the first to the third quartile, with the line marking the median. The whiskers extend to 1.5 times the box length.

energy error between 1.24 meV per atom and 4.67 meV per atom with mean error at 4.02 meV per atom. Finally, the unseen 1T' phase of MoS₂ returns an error range of -15.60 meV per atom to -4.34 meV per atom and a mean error of -8.06 meV per atom. These errors are all very competitive with that achieved by linear ML force fields constructed with a comparable range of parameters⁷⁵.

Next, we investigate the ability of our JLCDM to perform over systems never seen before. Our test is constructed for Al, for which we were able to build the best model, and consists in computing the total energy and forces of a series of 256-atom supercells taken from ref. ⁶⁴. This dataset contains 10 configurations corresponding to solid Al at 298 K and 10 configurations of both solid and liquid Al at its melting temperature of 933 K. The JLCDM used here is the same discussed before that produced the results from Fig. 2d–f, trained over 32-atom supercells for solid Al at 300 K. Table 2 summarizes our results. The error on the total energy and forces slightly increases when considering systems in the same conditions but different cell sizes, namely comparing the 32-atom and the 256-atom supercells for solid Al at 300 K and 298 K, respectively. In any case, the MAE remains below 1 meV per atom for the total energy and below 0.025 eV Å⁻¹ for the forces. As the structures tested become increasingly different from those used

Table 2. Performance of the JLCDM for Al, trained over 32-atom supercells at room temperature, against 256-atom supercells at various conditions.

# atoms	Condition	Total energy (meV per atom)		Forces (eV Å ⁻¹)	
		MAE	RMSE	MAE	RMSE
32	Solid (300 K)	0.046	0.054	0.007	0.009
256	Solid (298 K)	0.843	0.908	0.025	0.031
256	Solid (933 K)	6.976	7.526	0.068	0.862
256	Liquid (933 K)	35.062	36.498	0.164	0.203

The configurations for the 256-atom supercells are taken from refs. ^{64,80}, and the test error is computed over 10 samples for each different condition.

for training (data at 933 K) the error grows further, reaching 35.062 meV per atom and 0.164 eV Å⁻¹ in the liquid phase.

In order to put our results in perspective, neural network models ($\sim 10^6$ – 10^7 trainable weights) using the bispectrum components to describe the local environments reach a MAE of 123.29 meV per atom over the liquid phase, when trained on high-temperature solid structures only⁶⁴. This means that, on the same test, our JLCDM outperforms neural networks by a factor of four, despite consisting of only 120 trainable parameters and being trained on the 0.1% of the charge density points. The neural network error is then reduced to 13.04 meV per atom only when the training is performed on both high-temperature solids and liquids⁶⁴. Certainly, we could systematically improve the JLCDM by adding more distorted supercells in our training set or by including both solid and liquid configurations at 933 K. However, here, we wish to point out that the smooth description of the local environment allows us to achieve very competitive accuracy (35 meV per atom for liquid Al at 933 K) even for such a compact model.

DISCUSSION

Inspired by the recently developed Jacobi-Legendre potentials⁶⁶, we have designed a grid-based many-body linear expansion of the charge density, where the local external potential is described by Jacobi and Legendre polynomials. The method, combined with a charge-density targeted sampling strategy, produces highly accurate charge densities despite being constructed over an extremely limited number of trainable coefficients. We have demonstrated the efficacy of the JLCDM for diverse examples, namely a benzene molecule, solid and liquid Al, solid Mo and different phases of 2D MoS₂. In all cases, simple two-body JLCDMs accurately predict the charge density and can be transferred to different phases not originally included in the training set. For instance, training over the 1H and 1T phases of 2D MoS₂ is enough to predict the charge density of the 1T' phase, and so is the case for liquid Al, whose density can be constructed from a model trained over solid-state configurations at room temperature. The JLCDM-predicted densities can then be used to compute total energy and forces, achieving accuracy comparable to state-of-the-art machine learning force fields and, in some cases, even to fully converged DFT calculations.

As it stands, the methodology introduced here could be readily used in a diverse set of applications. If one is interested solely in energies and forces, learning the charge density probably will not be the optimal way to address the problem because of the computational overheads involved in many of the steps required for training and predicting the charge density over a fine grid. In that case, ML force fields can be a better solution, even though the

numerical effort to generate the training set needed to achieve DFT accuracy is typically rather extensive, much larger than that required to generate a JLCDM. In any case, a JLCDM strategy becomes essential when one targets density-related electronic quantities, which can be obtained only by DFT. For instance, one may need to evaluate the dipole moment (the Bader charges, the polarizability, etc.) along a set of molecular dynamics trajectories. In that case, a successful strategy may be to use a force field to generate the trajectories and the JLCDM method to evaluate the charge density and the associated properties. Furthermore, applications such as crystal structure prediction, phase diagram construction, reaction path search, and other computationally intensive tasks could be greatly accelerated by using JLCDM-predicted charge densities as the starting point of DFT calculations. Finally, the predicted charge density can be easily employed as the starting density for computationally expensive hybrid-functional calculations.

METHODS

DFT calculations and dataset generation

All single-point and ab initio molecular dynamics (AIMD) calculations are performed using density functional theory (DFT)^{1,2} as implemented in the Vienna ab initio simulation package (VASP)^{67,68}. Exchange and correlation interactions are treated by the generalized gradient approximation (GGA)³ with the Perdew–Burke–Ernzerhof (PBE)⁴ exchange and correlation functional. We use the projector augmented wave (PAW)⁷⁶ pseudopotentials. Single-point self-consistent calculations are performed with a 600 eV kinetic-energy cutoff for the plane-wave expansion, and the Brillouin zone is sampled over a k -point density of $12/\text{\AA}^{-3}$. AIMD runs are performed with a 2 fs time-step, and the Nosé–Hoover thermostat^{77–79} maintains the NVT ensemble. All AIMD runs are at least 4 ps long, and snapshots are taken from the simulation's last 3 ps. For benzene and 2D MoS₂ sufficient vacuum space, at least 15 Å is included in the non-periodic directions so to avoid spurious interaction between periodic images.

Benzene data generation. Data for benzene are extracted from the dataset available at <http://quantum-machine.org/datasets/62>. For the training set, we randomly select 30 snapshots from a MD run at 300 K and 400 K, available in the “benzene_300K-400K.tar.gz” file, and for the test set, 30 snapshots are randomly sampled from MD at 300 K, available in “benzene_300K-test.tar.gz”. The charge density for the selected snapshots is then calculated using VASP with the settings described above. Using 600 eV as the kinetic-energy cutoff for the plane-wave expansion, this results in the charge density being represented over a $180 \times 180 \times 180$ grid (5,832,000 grid points).

Al, Mo, and 2D MoS₂ data generation. For Al, Mo and MoS₂, we randomly extract snapshots from AIMD runs at 300 K. For Al, we use a $2 \times 2 \times 2$ conventional *fcc* supercell containing 32 atoms, while a $3 \times 3 \times 3$ conventional *bcc* supercell containing 54 atoms described Mo. A $3 \times 3 \times 1$ supercell is used for the 1H and 1T phases of MoS₂ (27 atoms), while for the 1T', we consider a $4 \times 2 \times 1$ supercell (48 atoms). For Al and Mo, we extract 10 snapshots for training and 10 for testing. For MoS₂, we extract 10 snapshots for each phase, with the 1H and 1T structures used for training and 1T' for testing. The charge densities are represented over a $140 \times 140 \times 140$ grid (2,744,000 grid points) for Al, a $160 \times 160 \times 160$ grid (4,096,000 grid points) for Mo, $160 \times 160 \times 300$ grid (7,680,000 grid points) for MoS₂ 1H and 1T, and $216 \times 192 \times 300$ grid (12,441,600 grid points) for 1T'-MoS₂.

In order to investigate the transferability of the JLCDM for Al, we use the snapshots reported in ref. ⁶⁴, as available in ⁸⁰. These Al are 256-atom Al supercells whose charge density has been recalculated with VASP. The energy cutoff for those is lowered to 360 eV

so as to match the same real-space grid used in ref. ⁶⁴, $200 \times 200 \times 200$ (8,000,000 grid points), and only the Γ -point is used to sample the BZ.

DFT calculations with fixed charge density

In order to use the ML charge density to compute total energies and forces, we use KS-DFT while keeping the charge density fixed and using the same settings as specified for the data generation. The ML charge density is kept constant at each step of an iterative diagonalization of the Kohn-Sham Hamiltonian. In particular, the Kohn-Sham eigenstates and eigenvalues are optimized during five steps with no updates to the charge density. A comparison for the Al test set at each step of the (non-) self-consistent cycle can be found in the SI.

While using PAW pseudopotentials, one is required to provide the augmentation on-site occupancies at the start of a calculation. For Al, we ignore one-centre correction terms evaluated on the radial support grid, a strategy that allows us to use the charge-density predictions for unknown structures or arbitrary sizes. For the other systems, we reuse the already known one-centre occupation DFT-computed terms together with our ML charge density to start the new calculations for configurations on the test set. In the future, the augmentation occupancies can also be learned with a similar scheme as designed here. This will allow the use of the ML charge density as a starting point for DFT calculations of any structure.

Model training, hyperparameter optimization, and timing

We fit the linear models by using singular value decomposition to find the pseudo-inverse of A solving the matrix equation, $A\hat{x} = \hat{b}$, for the coefficients \hat{x} . Training and inference are performed using the Ridge class (with $\alpha = 0$) from the scikit-learn library⁸¹.

Hyperparameter optimization is performed through Bayesian optimization using Gaussian Processes (`gp_minimize`), as implemented in the scikit-optimize library⁸². This is done solely on part of the training set. For the Al and Mo JLCDMs, 8 training snapshots are used for training and the remaining 2 are for validation. For benzene, 27 are used for training, and 3 for validation. On MoS₂, we take one training snapshot for each phase (1H and 1T) as the validation set, and the remaining training snapshots are used for training. The optimization targets the minimization of the mean absolute error (MAE). Table 3 shows the hyperparameters for each model.

An assessment of the time needed to train our model and to perform new predictions is provided in Table 4, where this is compared to the time needed to perform a fully converged self-consistent (SCF) calculation for the same system. This enables us to evaluate the time saving achieved over a single non-self-consistent calculation performed with our JLCDM-predicted charge density. Note that an estimate of the total computation time saving, which is inclusive of the effort needed to generate

Table 3. Optimized hyperparameters and corresponding feature size for each model generated.

System	Body	r_{cut}	n_{max}	l_{max}	r_{min}	α	β	# features
Benzene	1b	2.80	27	–	–0.78	7.00	0.00	1572
	2b	2.80	12	5	0.00	7.00	0.00	
Al	1b	4.08	15	–	–0.74	7.87	3.62	120
	2b	4.08	6	6	0.00	5.87	1.75	
Mo	1b	4.04	20	–	–1.09	4.02	5.46	812
	2b	4.04	12	11	0.00	–0.08	2.38	
2D MoS ₂	1b	4.76	18	–	–0.93	6.72	6.97	2346
	2b	4.76	11	10	0.00	5.07	2.69	

Table 4. Timing assessment of the different steps needed to train the JLCDM and to perform new predictions with it.

	Time (s) 32-atom Al cell (2,744,000 grid points)	Time (s) 256-atom Al cell (8,000,000 grid points)
Training		
Sample grid points with smart sampling (training and test)	9.401	–
Generate fingerprint vectors from atom positions (training and test)	25.424	–
Fitting linear model	3.449	–
Total training time	38.274	–
Inference		
Generate fingerprint vectors from atoms' positions	247.105	710.196
Construct density using JLCDM	0.085	0.088
DFT calculation using the JLCDM-predicted charge density	365.660	765.928
Total inference time	612.765	1,476.124
Fully converged SCF calculation	884.930	2,579.836
JLCDM saving	272.080 (30.75%)	1,103.624 (42.78%)

The time taken to create the training and test data (selecting the grid points using smart sampling and creating the fingerprints) refers to the 10 snapshots used for each data set partition. The fingerprint calculation is performed in the serial model (single CPU), while DFT calculations are performed with an MPI-parallel version of the VASP code run on 16 cores with the same parallelization settings. Note that the calculation for the 256-atom Al cell was performed with a JLCDM trained on the 32-atom cell so that there is only one training time.

the training set, strongly depends on the specific workflow one aims to follow since it scales with the total number of calculations to perform. Besides the DFT calculations, computing the JL fingerprints dominates the overall inference time. Importantly, we notice that the time taken for the inference of both a 32-atom and a 256-atom supercell is always inferior to the time of the full SCF-DFT calculation. As the size of the target system for inference increases, one sees a much more pronounced difference between the time needed for a JLCDM calculation and that for a full SCF-DFT one.

Grid sampling

The grid points included in the training set are selected by randomly sampling the real-space charge density according to a combination of uniform sampling and targeted sampling on the grid. Targeted sampling is performed by assigning to each grid point, \mathbf{r}_g , a probability P , given by a normal distribution of the inverse of the charge density at that grid point:

$$P(\mathbf{r}_g) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(1/n(\mathbf{r}_g))^2}{2\sigma^2}\right) \quad (10)$$

Targeted sampling is combined with uniform sampling across the simulation cell, composing the training data for each snapshot. The number of grid points sampled through targeted and uniform sampling is manually tuned to better sample the features of each example's charge density. The hyperparameter σ is also tuned manually for each example. However, these hyperparameters can readily be included in the automated

Table 5. Sampling hyperparameters and percentage of used data from the total data available.

System	σ	Uniform sampling (%)	Data used (%)
Benzene	90	50	0.10
Al	40	60	0.50
Mo	30	40	0.12
1T'-MoS ₂	40	60	0.05

The percentage of uniform sampling is retrieved out of the percentage of used data.

hyperparameter optimization routine, making it easy to address any molecule or solid-state system. Table 5 shows the parameters used for sampling and the percentage of the available grid point used to train the models. As shown in the Results section, our model requires a very modest data set size compared to other grid-based approaches in the literature while attaining accurate predictions.

In the SI we show the learning curve of the JLCD model for benzene with respect to the number of grid points, comparing our sampling technique with uniform sampling. Models trained with targeted-sampling data exhibit lower errors compared to those trained with uniform sampling, therefore, model accuracy is enhanced by using targeted sampling. In addition, the maximum absolute error for uniform sampling is significantly higher than that of targeted sampling, suggesting that the uniform sampling model has regions in space with poor prediction accuracy. We expect that as the data volume increases, the difference between uniform and targeted sampling diminishes. Nonetheless, we emphasize that our ability to use targeted sampling to develop highly accurate models with minimal data contributes to the efficiency of our workflow.

DATA AVAILABILITY

The data used to train and test the models (DFT charge density, structure files, and trained models) is available via Zenodo⁸³.

CODE AVAILABILITY

Scripts and related code for calculating the Jacobi-Legendre grid-based linear expansion are available at <https://github.com/StefanoSanvitoGroup/MLdensity>.

Received: 31 January 2023; Accepted: 17 May 2023;
Published online: 29 May 2023

REFERENCES

- Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
- Perdew, J. P. et al. Atoms, molecules, solids, and surfaces: applications of the generalized gradient approximation for exchange and correlation. *Phys. Rev. B* **46**, 6671–6687 (1992).
- Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- Burke, K. Perspective on density functional theory. *J. Chem. Phys.* **136**, 150901 (2012).
- Lignères, V. L. & Carter, E. A. in *An Introduction to Orbital-Free Density Functional Theory* 137–148 (Springer Netherlands, 2005).
- Nakata, A. et al. Large scale and linear scaling dft with the conquest code. *J. Chem. Phys.* **152**, 164112 (2020).
- Prentice, J. C. A. et al. The ONETEP linear-scaling density functional theory program. *J. Chem. Phys.* **152**, 174111 (2020).

9. Bartók, A. P. et al. Machine learning unifies the modeling of materials and molecules. *Sci. Adv.* **3**, e1701816 (2017).
10. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O. & Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
11. Morgani, D. & Jacobs, R. Opportunities and challenges for machine learning in materials science. *Annu. Rev. Mater. Res.* **50**, 71–103 (2020).
12. Zhuo, Y., Mansouri Tehrani, A. & Brgoch, J. Predicting the band gaps of inorganic solids by machine learning. *J. Phys. Chem. Lett.* **9**, 1668–1673 (2018).
13. Rajan, A. C. et al. Machine-learning-assisted accurate band gap predictions of functionalized MXene. *Chem. Mater.* **30**, 4031–4038 (2018).
14. Borlido, J. et al. Exchange-correlation functionals for band gaps of solids: benchmark, reparametrization and machine learning. *npj Comput. Mater.* **6**, 96 (2020).
15. Stanev, V. et al. Machine learning modeling of superconducting critical temperature. *npj Comput. Mater.* **4**, 29 (2018).
16. Xie, S. R., Stewart, G. R., Hamlin, J. J., Hirschfeld, P. J. & Hennig, R. G. Functional form of the superconducting critical temperature from machine learning. *Phys. Rev. B* **100**, 174513 (2019).
17. Nelson, J. & Sanvito, S. Predicting the curie temperature of ferromagnets using machine learning. *Phys. Rev. Mater.* **3**, 104405 (2019).
18. Xie, S. R. et al. Machine learning of superconducting critical temperature from Eliashberg theory. *npj Comput. Mater.* **8**, 14 (2022).
19. Zhang, J. et al. Machine learning prediction of superconducting critical temperature through the structural descriptor. *J. Phys. Chem. C* **126**, 8922–8927 (2022).
20. Bartel, C. J. et al. A critical examination of compound stability predictions from machine-learned formation energies. *npj Comput. Mater.* **6**, 97 (2020).
21. Claussen, N., Bernevig, B. A. & Regnault, N. Detection of topological materials with machine learning. *Phys. Rev. B* **101**, 245117 (2020).
22. Schleder, G. R., Focassio, B. & Fazio, A. Machine learning for materials discovery: two-dimensional topological insulators. *Appl. Phys. Rev.* **8**, 031409 (2021).
23. Chibani, S. & Coudert, F.-X. Machine learning approaches for the prediction of materials properties. *APL Mater.* **8**, 080701 (2020).
24. Schleder, G. R., Padilha, A. C. M., Acosta, C. M., Costa, M. & Fazio, A. From DFT to machine learning: recent approaches to materials science—a review. *J. Phys. Mater.* **2**, 032001 (2019).
25. Pickard, C. J. & Needs, R. J. Ab initio random structure searching. *J. Phys. Condens. Matter* **23**, 053201 (2011).
26. Podryabinkin, E. V., Tikhonov, E. V., Shapeev, A. V. & Oganov, A. R. Accelerating crystal structure prediction by machine-learning interatomic potentials with active learning. *Phys. Rev. B* **99**, 064114 (2019).
27. Tong, Q. et al. Combining machine learning potential and structure prediction for accelerated materials design and discovery. *J. Phys. Chem. Lett.* **11**, 8710–8720 (2020).
28. Wengert, S., Csányi, G., Reuter, K. & Margraf, J. T. Data-efficient machine learning for molecular crystal structure prediction. *Chem. Sci.* **12**, 4536–4546 (2021).
29. Deringer, V. L., Caro, M. A. & Csányi, G. A general-purpose machine-learning force field for bulk and nanostructured phosphorus. *Nat. Commun.* **11**, 5461 (2020).
30. Muhli, H. et al. Machine learning force fields based on local parametrization of dispersion interactions: application to the phase diagram of C60. *Phys. Rev. B* **104**, 054106 (2021).
31. Gubaev, K., Podryabinkin, E. V., Hart, G. L. & Shapeev, A. V. Accelerating high-throughput searches for new alloys with active learning of interatomic potentials. *Comput. Mater. Sci.* **156**, 148–156 (2019).
32. Rosenbrock, C. W. et al. Machine-learned interatomic potentials for alloys and alloy phase diagrams. *npj Comput. Mater.* **7**, 24 (2021).
33. Jinnouchi, R., Karsai, F. & Kresse, G. On-the-fly machine learning force field generation: application to melting points. *Phys. Rev. B* **100**, 014105 (2019).
34. Torres, A., Pedroza, L. S., Fernandez-Serra, M. & Rocha, A. R. Using neural network force fields to ascertain the quality of ab initio simulations of liquid water. *J. Phys. Chem. B* **125**, 10772–10778 (2021).
35. Jennings, P. C., Lysgaard, S., Hummelshøj, J. S., Vegge, T. & Bligaard, T. Genetic algorithms for computational materials discovery accelerated by machine learning. *npj Comput. Mater.* **5**, 46 (2019).
36. Cobelli, M., Cahalane, P. & Sanvito, S. Local inversion of the chemical environment representations. *Phys. Rev. B* **106**, 035402 (2022).
37. Choudhary, K., DeCost, B. & Tavazza, F. Machine learning with force-field-inspired descriptors for materials: fast screening and mapping energy landscape. *Phys. Rev. Mater.* **2**, 083801 (2018).
38. Tkatchenko, A. Machine learning for chemical discovery. *Nat. Commun.* **11**, 4125 (2020).
39. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
40. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
41. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
42. Huo, H. & Rupp, M. Unified representation of molecules and crystals for machine learning. *Mach. Learn. Sci. Technol.* **3**, 045017 (2022).
43. Zuo, Y. et al. Performance and cost assessment of machine learning interatomic potentials. *J. Phys. Chem. A* **124**, 731–745 (2020).
44. Himanen, L. et al. DScribe: library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* **247**, 106949 (2020).
45. Unke, O. T. et al. Machine learning force fields. *Chem. Rev.* **121**, 10142–10186 (2021).
46. Snyder, J. C., Rupp, M., Hansen, K., Müller, K.-R. & Burke, K. Finding density functionals with machine learning. *Phys. Rev. Lett.* **108**, 253002 (2012).
47. Nagai, R., Akashi, R. & Sugino, O. Completing density functional theory by machine learning hidden messages from molecules. *npj Comput. Mater.* **6**, 43 (2020).
48. Dick, S. & Fernandez-Serra, M. Machine learning accurate exchange and correlation functionals of the electronic density. *Nat. Commun.* **11**, 3509 (2020).
49. Li, L. et al. Kohn-Sham equations as regularizer: building prior knowledge into machine-learned physics. *Phys. Rev. Lett.* **126**, 036401 (2021).
50. Kirkpatrick, J. et al. Pushing the frontiers of density functionals by solving the fractional electron problem. *Science* **374**, 1385–1389 (2021).
51. Nelson, J., Tiwari, R. & Sanvito, S. Machine learning density functional theory for the Hubbard model. *Phys. Rev. B* **99**, 075132 (2019).
52. Snyder, J. C. et al. Orbital-free bond breaking via machine learning. *J. Chem. Phys.* **139**, 224104 (2013).
53. Meyer, R., Weichselbaum, M. & Hauser, A. W. Machine learning approaches toward orbital-free density functional theory: simultaneous training on the kinetic energy density functional and its functional derivative. *J. Chem. Theory Comput.* **16**, 5685–5694 (2020).
54. Alghadeer, M., Al-Aswad, A. & Alharbi, F. H. Highly accurate machine learning model for kinetic energy density functional. *Phys. Lett. A* **414**, 127621 (2021).
55. Ryczko, K., Wetzel, S. J., Melko, R. G. & Tamblin, I. Toward orbital-free density functional theory with small data sets and deep learning. *J. Chem. Theory Comput.* **18**, 1122–1128 (2022).
56. Pederson, R., Kalita, B. & Burke, K. Machine learning and density functional theory. *Nat. Rev. Phys.* **4**, 357–358 (2022).
57. Zhang, L. et al. Equivariant analytical mapping of first principles Hamiltonians to accurate and transferable materials models. *npj Comput. Mater.* **8**, 158 (2022).
58. Schütt, K. T., Gastegger, M., Tkatchenko, A., Müller, K.-R. & Maurer, R. J. Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions. *Nat. Commun.* **10**, 5024 (2019).
59. Hermann, J., Schätzle, Z. & Noé, F. Deep-neural-network solution of the electronic Schrödinger equation. *Nat. Chem.* **12**, 891–897 (2020).
60. Grisafi, A. et al. Transferable machine-learning model of the electron density. *ACS Cent. Sci.* **5**, 57–64 (2019).
61. Lewis, A. M., Grisafi, A., Ceriotti, M. & Rossi, M. Learning electron densities in the condensed phase. *J. Chem. Theory Comput.* **17**, 7203–7214 (2021).
62. Brockherde, F. et al. Bypassing the Kohn-Sham equations with machine learning. *Nat. Commun.* **8**, 872 (2017).
63. Chandrasekaran, A. et al. Solving the electronic structure problem with machine learning. *npj Comput. Mater.* **5**, 22 (2019).
64. Ellis, J. A. et al. Accelerating finite-temperature Kohn-Sham density functional theory with deep neural networks. *Phys. Rev. B* **104**, 035120 (2021).
65. Rackers, J. A., Tecot, L., Geiger, M. & Smidt, T. E. A recipe for cracking the quantum scaling limit with machine learned electron densities. *Mach. Learn. Sci. Technol.* **4**, 015027 (2023).
66. Domina, M., Patil, U., Cobelli, M. & Sanvito, S. The Jacobi-Legendre potential. Preprint at <https://arxiv.org/abs/2208.10292> (2022).
67. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186 (1996).
68. Kresse, G. & Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**, 15–50 (1996).
69. Kohn, W. & Yaniv, A. Locality principle in wave mechanics. *Proc. Natl Acad. Sci. USA* **75**, 5270–5272 (1978).
70. Pozdnyakov, S. N. et al. Incompleteness of atomic structure representations. *Phys. Rev. Lett.* **125**, 166001 (2020).
71. Eda, G. et al. Coherent atomic and electronic heterostructures of single-layer MoS₂. *ACS Nano* **6**, 7311–7317 (2012).
72. Yang, D., Sandoval, S. J., Divigalpitaya, W. M. R., Irwin, J. C. & Frindt, R. F. Structure of single-molecular-layer mos₂. *Phys. Rev. B* **43**, 12053–12056 (1991).
73. Mattheiss, L. F. Band structures of transition-metal-dichalcogenide layer compounds. *Phys. Rev. B* **8**, 3719–3740 (1973).
74. Qian, X., Liu, J., Fu, L. & Li, J. Quantum spin Hall effect in two-dimensional transition metal dichalcogenides. *Science* **346**, 1344–1347 (2014).

75. Lunghi, A. & Sanvito, S. A unified picture of the covalent bond within quantum-accurate force fields: from organic molecules to metallic complexes' reactivity. *Sci. Adv.* **5**, eaaw2210 (2019).
76. Kresse, G. & Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **59**, 1758–1775 (1999).
77. Nosé, S. A unified formulation of the constant temperature molecular dynamics methods. *J. Chem. Phys.* **81**, 511–519 (1984).
78. Nosé, S. A molecular dynamics method for simulations in the canonical ensemble. *Mol. Phys.* **100**, 191–198 (1984).
79. Hoover, W. G. Canonical dynamics: equilibrium phase-space distributions. *Phys. Rev. A* **31**, 1695–1697 (1985).
80. Ellis, J. A. et al. LDOS/SNAP data for MALA: aluminium at 298K and 933K. *rodare* <https://doi.org/10.14278/rodare.1107> (2021).
81. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
82. Head, T., Kumar, M., Nahrstaedt, H., Louppe, G. & Shcherbatyi, I. scikit-optimize. zenodo <https://doi.org/10.5281/zenodo.5565057> (2021).
83. Focassio, B., Domina, M., Patil, U., Fazzio, A. & Sanvito, S. Linear Jacobi-Legendre expansion of the charge density for machine learning-accelerated electronic structure calculations. zenodo <https://doi.org/10.5281/zenodo.7922012> (2023).

ACKNOWLEDGEMENTS

This work was supported by São Paulo Research Foundation (FAPESP) (Grant nos. 2021/12204-6, 2019/04527-0, and 2017/02317-2) and by the Irish Research Council Advanced Laureate Award (IRCLA/2019/127). We acknowledge the DJEI/DES/SFI/HEA Irish Centre for High-End Computing (ICHEC) and Trinity Centre for High Performance Computing (TCHPC) for the provision of computational resources. We acknowledge support from ICHEC via the academic flagship programme (Project Number - EuroCC-AF-3). We acknowledge NVIDIA Academic Hardware Grant Program for providing graphics processing units.

AUTHOR CONTRIBUTIONS

S.S. conceived the idea of a machine learning model for the charge density as starting guess for DFT calculations. M.D. developed the Jacobi–Legendre representation of the grid points and the many-body linear expansion of the charge density. B.F. and U.P. implemented the grid-based JL representation and linear model. B.F. performed

all DFT calculations, ML training, and testing and implemented related code for results analysis. S.S. and A.F. supervised the work. All authors contributed to discussions, writing, and revision of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-023-01053-0>.

Correspondence and requests for materials should be addressed to Bruno Focassio or Stefano Sanvito.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023