By-passing the Kohn-Sham equations with machine learning

Felix Brockherde,^{1,2} Leslie Vogt,³ Li Li,⁴ Mark E. Tuckerman,^{3,5,6} Kieron Burke,^{7,4,*} and Klaus-Robert Müller^{1,8,9,*}

¹Machine Learning Group, Technische Universität Berlin, Marchstr. 23, 10587 Berlin, Germany

² Max-Planck-Institut für Mikrostrukturphysik, Weinberg 2, 06120 Halle, Germany

³Department of Chemistry, New York University, New York, NY 10003, USA

⁴Departments of Physics and Astronomy, University of California, Irvine, CA 92697, USA

⁵Courant Institute of Mathematical Science, New York University, New York, NY 10003, USA

⁶NYU-ECNU Center for Computational Chemistry at NYU Shanghai,

3663 Zhongshan Road North, Shanghai 200062, China

⁷Departments of Chemistry, University of California, Irvine, CA 92697, USA

⁸Department of Brain and Cognitive Engineering, Korea University,

Anam-dong, Seongbuk-gu, Seoul 136–713, Republic of Korea

⁹Max Planck Institute for Informatics, Stuhlsatzenhausweg, 66123 Saarbrücken, Germany

(Dated: June 15, 2017)

Last year, at least 30,000 scientific papers used the Kohn-Sham scheme of density functional theory to solve electronic structure problems in a wide variety of scientific fields, ranging from materials science to biochemistry to astrophysics. Machine learning holds the promise of learning the kinetic energy functional via examples, by-passing the need to solve the Kohn-Sham equations. This should yield substantial savings in computer time, allowing either larger systems or longer time-scales to be tackled, but attempts to machine-learn this functional have been limited by the need to find its derivative. The present work overcomes this difficulty by directly learning the density-potential and energy-density maps for test systems and various molecules. Both improved accuracy and lower computational cost with this method are demonstrated by reproducing DFT energies for a range of molecular geometries generated during molecular dynamics simulations. Moreover, the methodology could be applied directly to quantum chemical calculations, allowing construction of density functionals of quantum-chemical accuracy.

INTRODUCTION

Kohn-Sham density functional theory[1] is now enormously popular as an electronic structure method in a wide variety of fields[2]. Useful accuracy is achieved with standard exchange-correlation approximations, such as generalized gradient approximations[3] and hybrids[4]. Such calculations are playing a key role in the materials genome initiative[5], at least for weakly correlated materials[6].

There has also been a recent spike of interest in applying machine learning (ML) methods in the physical sciences[7–11]. The majority of these applications involve predicting properties of molecules or materials from large databases of KS-DFT calculations[12–15]. A few applications involve finding potential energy surfaces within MD simulations[16–19]. Fewer still have focussed on finding the functionals of DFT as a method of performing KS electronic structure calculations without solving the KS equations[20–23]. If such attempts could be made practical, the possible speed-up in repeated DFT calculations of similar species, such as occur in ab initio MD simulations, is enormous.

A key difficulty has been the need to extract the functional derivative of the non-interacting kinetic energy. The non-interacting kinetic energy functional $T_{\rm s}[n]$ of the density n is used in *two* distinct ways in a KS calculation[1], as illustrated in Fig. 1: (i) its functional derivative is used in the Euler equation which is solved

in the self-consistent cycle and (ii) when self-consistency is reached, the ground-state energy of the system is calculated by E[n], an Orbital-Free (OF) mapping. The solution of the KS equations performs both tasks exactly. Early results on simple model systems showed that machine learning could provide highly accurate values for $T_{\rm s}[n]$ with only modest amounts of training[20], but that the corresponding functional derivatives are too noisy to yield sufficiently accurate results to (i). Subsequent schemes overcome this difficulty in various ways, but typically lose a factor of 10 or more in accuracy[22], and their computational cost can increase dramatically with system complexity.

Here we present an alternative ML approach, in which we replaced the Euler equation by directly learning the Hohenberg-Kohn (HK) map $v(\mathbf{r}) \rightarrow n(\mathbf{r})$ (red line in Fig. 1a) from the one-body potential of the system of interest to the interacting ground-state density, i.e. we establish an ML-HK map. We show that this map can be learned at a much more modest cost than either previous ML approaches to find the functional and its derivative (ML-OF) or direct attempts to model the energy as a functional of $v(\mathbf{r})$ (ML-KS). Furthermore we show that it can immediately be applied to molecular calculations, by calculating the energies of small molecules over a range of conformers. Moreover, since we have already implemented this approach with a standard quantum chemical code (Quantum Espresso[24]) using a standard DFT approximation (PBE), this can now be tried on much larger scales.

The ML-HK map reflects the underlying computational approach used to generate a particular electron density, but is not restricted to any given electronic structure method. Many molecular properties, not only the energy, are dependent on the electron density, making the ML-HK map more versatile than a direct ML-KS mapping. We also establish that densities can be learned with sufficient accuracy to distinguish between different DFT functionals, providing a route to future functional development by generating precise densities for a range of molecules and conformations.

RESULTS

We will first outline theoretical results, most prominently the ML-HK map, and then illustrate the approach with simulations of 1-D systems and 3-D molecules.

ML-Hohenberg-Kohn map

Previous results show that for an ML-OF approach, the accuracy of ML KS kinetic energy models $T_{\rm s}^{\rm ML}[n]$ improve rapidly with the amount of data. But minimizing the total energy via gradient descent requires the calculation of the gradient of the kinetic energy model $T_{\rm s}^{\rm ML}$ (see Fig. 1). Calculating this gradient is challenging. Due to the data driven nature of, e.g., kernel models, the machine-learned kinetic energy functional has no information in directions that point outside the data manifold[25]. This heavily influences the gradient to an extent that it becomes unusable without further processing[20]. There have been several suggestions to remedy this problem but all of them share a significant loss in accuracy compared to $T_{\rm s}[n][21, 22, 26]$.

However, we propose an interesting alternative to gradients and the ML-OF approach. Recently, it has been shown that the Hohenberg-Kohn map for the density as a functional of the potential can be approximated extremely accurately using semiclassical expressions[27]. Such expressions do not require the solution of any differential equation, and become more accurate as the number of particles increases. Errors can be negligible even for just 2 distinct occupied orbitals.

Inspired by this success, we suggest to circumvent the kinetic energy gradient and directly train a multivariate machine learning model. We name this the ML-Hohenberg-Kohn (ML-HK) map:

$$n^{\rm ML}[v](x) = \sum_{i=1}^{M} \beta_i(x)k(v, v_i).$$
 (1)

Here, each density grid point is associated with a group of model weights β . Training requires solving an optimization problem for each density grid point. While this is possible in 1-D, it rapidly becomes intractable in 3-D, since the number of grid points grows cubically.

The use of a basis representation for the densities, as in

$$n^{\rm ML}[v](x) = \sum_{l=1}^{L} u^{(l)}[v]\phi_l(x), \qquad (2)$$

renders the problem tractable even for 3-D. A machine learning model that predicts the basis function coefficients $u^{(l)}[v]$ instead of the grid points is then formulated.

Predicting the basis function coefficients not only makes the machine learning model efficient and allows the extension of the approach to 3-D but also permits regularization, e.g. to smooth the predicted densities by removing the high frequency basis functions for example, or to further regularize the machine learning model complexity for specific basis functions.

For orthogonal basis functions, the machine learning model reduces to several independent regression models and admits an analytical solution analogous to Kernel Ridge Regression (see supplement Eq. 5):

$$\boldsymbol{\beta}^{(l)} = \left(\mathbf{K}_{\sigma^{(l)}} + \lambda^{(l)} \mathbf{I} \right)^{-1} \mathbf{u}^{(l)}, \quad l = 1, \dots, L. \quad (3)$$

Here, for each basis function coefficient, $\lambda^{(l)}$ are regularization parameters and $K_{\sigma^{(l)}}$ is a Gaussian kernel with kernel width $\sigma^{(l)}$. The $\lambda^{(l)}$ and $\sigma^{(l)}$ can be chosen individually for each basis function via independent crossvalidation (see [12, 28]). This ML-HK model avoids prior gradient descent procedures and with it the necessity to "de-noise" the gradients. Due to the independence of Eq. 3 for each l, the solution scales nicely.

Functional and Density driven error

How can the performance of the ML-HK map be measured? It has recently been shown how to separate out the effect of the error in the functional F and the error in the density $n(\mathbf{r})$ on the resulting error in the total energy of any approximate, self-consistent DFT calculation[29]. Let \tilde{F} be an approximation of the many body functional F, and $\tilde{n}(\mathbf{r})$ the approximate ground-state density when \tilde{F} is used in the Euler equation. Defining $\tilde{E}[n] = \tilde{F}[n] + \int d^3r n(\mathbf{r}) \mathbf{v}(\mathbf{r})$ yields

$$\Delta E = \tilde{E}[\tilde{n}] - E[n] = \Delta E_F + \Delta E_D \tag{4}$$

where $\Delta E_F = \tilde{F}[n] - F[n]$ is the functional-driven error, while $\Delta E_D = \tilde{E}[\tilde{n}] - \tilde{E}[n]$ is the density-driven error.



Figure 1. **a.** Mappings used in this paper. The bottom arrow represents E[v], a conventional electronic structure calculation, i.e., KS-DFT. The ground state energy is found by solving KS equations given the external potential, v. E[n] is the total energy density functional. The red arrow is the HK map n[v] from external potential to its ground state density. **b** top. How the energy error depends on M for ML-OF and ML-HK with different basis sets for the 1-D problem. **b** bottom. Errors of the PBE energies (relative to exact values) and the ML maps (relative to PBE) as a function of interatomic spacing, R, for H_2 with M = 7. **c**. How our Machine Learning Hohenberg-Kohn (ML-HK) map makes predictions. The molecular geometry is represented by Gaussians; many independent Kernel Ridge Regression models predict each basis coefficient of the density. We analyze the performance of data-driven (ML) and common physical basis representations for the electron density.

In most DFT calculations, ΔE is dominated by ΔE_F . The standard DFT approximations can, in some specific cases, produce abnormally large density errors that dominate the total error. In such situations, using a more accurate density can greatly improve the result [29–31]. We will use these definitions to measure the accuracy of the ML-HK map.

1-D potentials

The following results demonstrate how much more accurate ML is when applied directly to the HK map. The box problem originally introduced in Snyder *et al.* [20] is used to illustrate the principle. Random potentials consisting of three Gaussian dips were generated inside a hard-wall box of length 1 (atomic units), and the Schrödinger equation for one electron was solved extremely precisely. Up to 200 cases were used to train an ML model $T_{\rm s}^{\rm ML}[n]$ for the non-interacting kinetic energy functional $T_{\rm s}[n]$ via Kernel Ridge Regression (for details, see supplement).

To measure the accuracy of an approximate HK map, the analysis of the previous section is applied to the KS DFT problem. Here F is just T_s , the non-interacting kinetic energy, and

$$\Delta E_F = \tilde{T}_{\rm s}[n] - T_{\rm s}[n], \qquad (5)$$

i.e., the error made in an approximate functional on the

	ML-OF					ML-HK (grid)					ML-HK (other)					
	ΔE		ΔE_F		ΔE_D		ΔE		ΔE_D		ΔE_D^{ML}		ΔE_D (Fourier)		ΔE_D (KPCA)	
M	MAE	max	MAE	max	MAE	max	MAE	max	MAE	max	MAE	max	MAE	max	MAE	max
20	7.7	47	7.7	60	8.8	87	3.5	27	0.76	8.9	9.7	70	0.58	8	0.15	2.9
50	1.6	30	1.3	7.3	1.4	31	1.2	7.1	0.079	0.92	0.27	2.4	0.078	0.91	0.011	0.17
100	0.74	17	0.2	2.6	0.75	17	0.19	2.1	0.027	0.43	0.18	2.4	0.031	0.42	0.0012	0.028
200	0.17	2.9	0.039	0.6	0.17	2.9	0.042	0.59	0.0065	0.15	0.02	0.46	0.017	0.14	0.00055	0.015

Table I. Energy errors in kcal/mol for the 1-D data set for various M, the number of training points. For definitions, see text.

exact density. Table I on the left gives the errors made by ML-OF for the total energy, and its different components, when the density is found from the functional derivative. This method works by following a gradient descent of the total energy functional based on the gradient of the ML model $T_{\rm s}^{\rm ML}$,

$$n^{(j+1)} = n^{(j)} - \epsilon P\left(n^{(j)}\right) \frac{\delta}{\delta n} E^{\mathrm{ML}}(n^{(j)}), \qquad (6)$$

where ϵ is a small number and $P(n^{(j)})$ is a localized PCA projection to de-noise the gradient. Here and for all further 1-D results we use

$$E^{\rm ML}[n] = T_{\rm s}^{\rm ML}[n] + \int \mathrm{d}x \, n(x) \, v(x).$$
 (7)

The density-driven contribution to the error ΔE_D , which we calculate exactly here using the von Weizsäcker kinetic energy[32] is always comparable to, or greater than, the functional-driven error ΔE_F , due to the poor quality of the ML functional derivative[20]. The calculation is abnormal, and can be greatly improved by using a more accurate density from a finer grid. As the number of training points M grows, the error becomes completely dominated by the error in the density. This shows that the largest source of error is in using the ML approximation of T_s to find the density by solving the Euler equation.

The next set of columns analyzes the ML-HK approach, using a grid basis. The left-most of these columns shows the energy error we obtain by utilizing the ML-HK map:

$$\Delta E = |E^{\mathrm{ML}}[n^{\mathrm{ML}}[v]] - E|. \tag{8}$$

Note that both ML models, $T_{\rm s}^{\rm ML}$ and $n^{\rm ML}$, have been trained using the same set of M training points.

The ML-HK approach is always more accurate than ML-OF, and its relative performance improves as M increases. The next column reports the density-driven error ΔE_D which is an order-of-magnitude smaller than for ML-OF. Lastly, we list an estimate to the density-driven error

$$\Delta E_D^{\mathrm{ML}} = |E^{\mathrm{ML}}[n^{\mathrm{ML}}[v]] - E^{\mathrm{ML}}[n]|, \qquad (9)$$

which uses the ML model $T_{\rm s}^{\rm ML}$ for the kinetic energy functional in 1-D. This proxy is generally a considerable overestimate (a factor of 3 too large), so that the true ΔE_D is always significantly smaller. We use it in subsequent calculations (where we cannot calculate $T_{\rm s}^{\rm ML}$) to (over-)estimate the energy error due to the HK-ML map.

The last set of columns are density-driven errors for other basis sets. Three variants of the ML-HK map were tested. First, direct prediction of the grid coefficients: In this case, $\mathbf{u}_i^{(l)} = n_i(x_l), \ l = 1, \dots, G$. 500 grid points were used, as in Snyder et al. [20]. This variant is tested in 1-D only; in 3-D the high dimensionality will be prohibitive. Second, a common Fourier basis is tested. The density can be transformed efficiently via the discrete Fourier transform, using 200 Fourier basis functions in total. In 3-D these basis functions correspond to plane waves. The back-projection $\mathbf{u} \mapsto n$ to input space is simple, but although the basis functions are physically motivated, they are very general and not specifically tailored to density functions. The performance is almost identical to the grid on average, although maximum errors are much less. For M = 20, the error that originates from the basis representation starts to dominate. This is a motivation for exploring, third, a Kernel PCA (KPCA) basis[33]. KPCA[34] is a popular generalization of PCA that yields basis functions that maximize variance in a higher dimensional feature space. The KPCA basis functions are data-driven and computing them requires an eigen-decomposition of the Kernel matrix. Good results are achieved with only 25 KPCA basis functions. The KPCA approach gives better results because it can take the non-linear structure in the density space into account. However, it introduces the pre-image problem: It is not trivial to project the densities from KPCA space back to their original (grid) space (see supplement). It is thus not immediately applicable to 3-D applications.

Molecules

We next apply the ML-HK approach to predict electron densities and energies for a series of small molecules. We test the ML models on KS-DFT results obtained using the PBE exchange-correlation functional[35] and atomic pseudopotentials with the projector augmented wave (PAW) method[36, 37] in the Quantum ESPRESSO code.[38] Since the ML-OF approach applied in the previous section becomes prohibitively expensive in 3-D due to the poor convergence of the gradient descent procedure, we compare the ML-HK map to the ML-KS approach. This approach models the energy directly as a functional of $v(\mathbf{r})$, i.e. it trains a model

$$E^{\mathrm{ML}}[v] = \sum_{i=1}^{M} \alpha_i k(v_i, v) \tag{10}$$

via KRR (for details, see supplement).

We also apply the ML-HK map with Fourier basis functions. Instead of a $T_{\rm s}^{\rm ML}[n]$ model, we learn an $E^{\rm ML}[n]$ model

$$E^{\mathrm{ML}}[n] = \sum_{i=1}^{M} \alpha_i k(n_i, n) \tag{11}$$

which avoids implementing the potential energy and exchange-correlation functionals.

Both approaches require the characterization of the Hamiltonian by its external potential. The external (Coulomb) potential diverges for the 3-D molecules and is therefore not a good feature to measure the distance in ML. Instead, we use an artificial *Gaussians potential* in the form of

$$v(\mathbf{r}) = \sum_{\alpha=1}^{N^a} Z_\alpha \exp\left(\frac{-\|\mathbf{r} - R_\alpha\|^2}{2\gamma^2}\right)$$
(12)

where R_{α} are the positions and Z_{α} are the nuclear charges of the N^a atoms. The Gaussians potential is used for the ML representation only. The width γ is a hyper-parameter of the algorithm. The choice is arbitrary but can be cross-validated. We find good results with $\gamma = 0.2$ Å. The idea of using Gaussians to represent the external potential has been used previously[39]. The Gaussians potential is discretized on a coarse grid with grid spacing $\Delta = 0.08$ Å. To use the discretized potential in the Gaussian kernel, we flatten it into a vector form and thus use a tensor Frobenius norm.

Our first molecular prototype is H_2 , with the only degree of freedom, R, denoting the distance between the H atoms. A dataset of 150 geometries is created by varying R between 0.5 and 1.5 Å (sampled uniformly). A randomly chosen subset of 50 geometries are designated as the *test set* and are unseen by the ML algorithms. These geometries are used to measure the *out-of-sample* error reported below.

The remaining 100 geometries make up the grand training set. To evaluate the performance of the ML-KS map and the ML-HK map, subsets of varying sizes M are chosen out of the grand training set to train the $E^{\rm ML}[v]$ and $n^{\rm ML}[v]$ models, respectively. Because the required training subsets are so small, careful selection of a subset that covers the complete range of R is necessary. This is accomplished by selecting the M training points out of the grand training set so that the R values are nearly equally spaced (see supplement for details).

For practical applications, it is not necessary to run DFT calculations for the complete grand training set, only for the M selected training points. Strategies for sampling the conformer space and selecting the grand training set for molecules with more degrees of freedom are explained for H_2O and MD simulations later on.

The performance of the ML-KS map and ML-HK map is compared by evaluating $E^{\rm ML}[v]$ that maps from the Gaussians potential to total energy and the combination of $n^{\rm ML}[v]$ that maps from Gaussians potential to the ground-state density in a three-dimensional Fourier basis representation (l = 25) and $E^{\rm ML}[n]$ that maps from density to total energy. The prediction errors are listed in Table II.

The MAE of the energy evaluated using the ML-HK map is significantly smaller than that of the ML-KS map. This indicates that even for a 3-D system, learning the potential-density relationship via the HK map is much easier than directly learning the potential-energy relationship via the KS map.

Fig. 1b shows the errors made by the ML-KS and the ML-HK maps. The error of the ML-HK map is smoother than the ML-KS error and is much smaller, even for the most problematic region when R is smaller than the equilibrium bond distance of $R_0 = 0.74$ Å. The MAE that is introduced by the PBE approximation on the H₂ dataset is 2.3 kcal/mol (compared to exact CI calculations), i.e., well above the errors of the ML model and verifies that the error introduced by the ML-HK map is negligible for a DFT calculation.

The next molecular example is H_2O , parametrized with three degrees of freedom: two bond lengths and a bond angle. To create a conformer dataset, the optimized structure ($R_0 = 0.97$ Å, $\theta_0 = 104.2^\circ$ using PBE) is taken as a starting point. A total of 350 geometries are then generated by changing each bond length by a uniformly sampled value between ± 0.075 Å and varying the angle θ between ± 8.59 degrees (± 0.15 rad) away from θ_0 (see supplement Fig. 1 for a visualization of the sampled range). For this molecule, the out-of-sample test set again comprises a random subset of 50 geometries, with the remaining 300 geometries used as the grand training set. Because there are now three parameters, it is

			ML-	KS		ML-HK							
		ΔE		ΔR_o	$\Delta \theta_0$	ΔE		ΔE_D^{ML}		ΔR_o	$\Delta \theta_0$		
Molecule	M	MAE	max			MAE	max	MAE	max				
	5	1.3	4.3	2.2		0.70	2.9	0.18	0.54	1.1			
H_2	7	0.37	1.4	0.23		0.17	0.73	0.054	0.16	0.19			
	10	0.080	0.41	0.23	_	0.019	0.11	0.017	0.086	0.073			
	5	1.4	5.0	2.1	2.2	1.1	4.9	0.056	0.17	2.3	3.8		
	10	0.27	0.93	0.63	1.9	0.12	0.39	0.099	0.59	0.12	0.38		
Н.О	15	0.12	0.47	0.19	0.41	0.043	0.25	0.029	0.14	0.064	0.23		
1120	20	0.015	0.064	0.043	0.16	0.0091	0.060	0.011	0.058	0.024	0.066		

Table II. Prediction errors on H_2 and H_2O with increasing number of training points M for the ML-KS and ML-HK approaches. In addition, the estimated density-driven contribution to the error for the ML-HK approach (Eq. 9) is given. Energies in kcal/mol, bond-lengths in pm, and angles in degrees.

more difficult to select equidistant samples for the training subset of M data points. We therefore use a K-means approach to find M clusters and select the grand training set geometry closest to each cluster's center for the training subset (see supplement for details).



Figure 2. **Top.** Distribution of energy errors against PBE on the H_2O dataset for ML-KS and ML-HK. The errors are plotted on a symmetric log scale with linear threshold of 0.01, using nearest neighbor interpolation from a grid scan for coloring. Black dots mark the test set geometries with averaged bond lengths. **Bottom left.** Comparison of the PBE errors made by ML-HK and ML-KS on the test set geometries. **Bottom right.** Energy landscape of the ML-HK map for symmetric geometries (R versus θ). All models trained on M = 15 training points. Energies and errors in kcal/mol. A black cross marks the PBE equilibrium position.

Models are trained as for H_2 . The results are given in Table II. As expected, the increase in degrees of freedom for H_2O compared to H_2 requires a larger training set size

M. However, even for the more complicated molecule, the ML-HK map is consistently more precise than the ML-KS map, and provides an improved potential energy surface, as shown in Fig. 2. With an MAE of 1.2 kcal/mol for PBE energies relative to CCSD(T) calculations for this data set, we again show that ML does not introduce a new significant source of error.

The ML maps can also be used to find the minimum energy configuration. The total energy is minimized as the geometry varies with respect to both bond lengths and angles. For optimization, we use Powell's method[40], which requires a starting point and an evaluation function to be minimized. For the H_2O case, the search is restricted to symmetric configurations, with a random symmetric geometry used as the starting point. Results are reported in Table II. The optimizations consistently converge to the correct minima regardless of starting point, consistent with the maps being convex, i.e., the potential energy curves are sufficiently smooth as to avoid introducing artificial local minima.

For larger molecules, generating random conformers that sample the full configurational space becomes difficult. Therefore, we next demonstrate that molecular dynamics (MD) using a classical force field can also be used to create the grand training set. As an example, we use benzene (C_6H_6) with only small fluctuations in atomic positions out of the molecular plane. Appropriate conformers are generated via isothermal MD simulations at 300 K, 350 K, and 400 K using the General Amber Force Field (GAFF)[41] in the PINY_MD package[42]. Saving snapshots from the MD trajectories generates a large set of geometries that are sampled using the K-means approach to obtain 2,000 representative points for the grand training set. Training $n^{ML}[v]$ and $E^{ML}[n]$ is performed as above by running DFT calculations on M = 2000 points. We find that the ML error is reduced by creating the training set from trajectories at both the target temperature and a higher temperature to increase the representation of more distorted geometries. The final ML model is tested on 200 conformational snapshots taken from an independent MD trajectory at 300 K (see Fig. 3a). The MAE of the ML-HK map for this data set using training geometries from 300 K and 350 K trajectories is only 0.37 kcal/mol for an energy range that spans more than 10 kcal/mol (see Table III).

For benzene, we further quantify the precision of the ML-HK map in reproducing PBE densities. In Fig. 4, it is clear that the errors in the Fourier basis representation are larger than the errors introduced by the ML-HK map by two orders of magnitude. Furthermore, the ML-HK errors in density (as evaluated on a grid in the molecular plane of benzene) are also considerably smaller than the difference in density between density functionals (PBE versus LDA[43]). This result verifies that the ML-HK map is specific to the density used to train the model and should be able to differentiate between densities generated with other electronic structure approaches.

Ethane (C_2H_6) , with a small energy barrier for the relative rotation of the methyl groups, is also evaluated in the same way. Using geometries sampled using K-means from 300 K and 350 K classical trajectories, the ML-HK model reproduces the energy of conformers with a MAE of 0.23 kcal/mol for an independent MD trajectory at 300 K (Fig. 3b). This test set includes conformers from the sparsely-sampled eclipsed configuration (see supplement Fig. 3). Using points from a 400 K trajectory improves the ML-HK map due to the increased probability of higher energy rotamers in the training set (see Table III). The training set could also be constructed by including explicit rotational conformers, as is common for fitting classical force field parameters[41]. In either case, generating appropriate conformers for training via computationally cheap classical MD significantly decreases the cost of the ML-HK approach.

As additional proof of the versatility of the ML-HK map, we show that this approach is also able to interpolate energies for proton transfer in the enol form of malonaldehyde $(C_3H_4O_2)$. This molecule is a well-known example of intramolecular proton transfer, and our previous AIMD and ab initio path integral studies [44] found classical and quantum free energy barrier values of 3.5 and 1.6 kcal/mol, respectively, from gradient-corrected DFT. In this work, classical MD trajectories are run for each tautomer separately, with a fixed bonding scheme, then combined for K-means sampling to create the grand training set. The training set also includes an artificially constructed geometry that is the average of tautomer atomic positions. For the test set, we use snapshots from a computationally expensive Born-Oppenheimer ab initio MD trajectory at 300 K. Fig. 5a shows that the ML-HK map is able to predict DFT energies during a proton transfer event (MAE of 0.27 kcal/mol) despite being trained on classical geometries that did not include these intermediate points.

The ML-HK map can also be used to *generate* a stable MD trajectory for malonaldehyde at 300 K (see Fig. 5b). In principle, analytic gradients could be obtained for each timestep, but for this first proof-of-concept trajec-

	Training	Δ	E	ΔE_D^{ML}		
Molecule	trajectories	MAE	max	MAE	max	
	300K	0.42	1.7	0.32	1.5	
Benzene	300K + 350K	0.37	1.8	0.28	1.5	
	300K + 400K	0.47	2.3	0.30	1.8	
	300K	0.20	1.5	0.17	1.3	
Ethane	300K + 350K	0.23	1.4	0.19	1.1	
	300K + 400K	0.14	1.7	0.098	0.62	
Malonaldehyde	300K + 350K	0.27	1.2	0.21	0.74	

Table III. Energy and density-driven errors (kcal/mol) of the ML-HK approach on the MD datasets for different training trajectory combinations.

tory, a finite-difference approach was used to determine atomic forces. The ML-HK-generated trajectory samples the same molecular configurations as the ab inito simulation (see Fig. 6), with mean absolute energy errors of 0.77 kcal/mol, but it typically underestimates the energy for out-of-plane molecular fluctuations at the extremes of the classical training set (maximum error of 5.7 kcal/mol). Even with underestimated energy values, the atomic forces are sufficiently large to return the molecule to the equilibrium configuration, resulting in a stable and long trajectory. The new set of coordinates could be further sampled to expand the training set in a self-consistent manner. Using iterative ML-HK-generated MD trajectories would eliminate the need to run computationally expensive MD simulations with DFT and would provide an iterative approach to reduce the energy errors for conformations not included in the classical training set.

DISCUSSION

For several decades, density functional theory has been a cross-disciplinary area between theoretical physics, chemistry, and materials sciences. The methods of each field cross-fertilize advances in other fields. This has led to its enormous popularity and widespread success, despite its well-known limitations in both accuracy and the systems and properties to which it can be applied.

The present work makes a key step forward toward adding an entirely new ingredient to this mix, namely the construction of functionals via machine learning. While previous work showed proofs of principle in 1-D, this is the first demonstration in 3-D, using real molecules and production-level codes. We also demonstrate that molecular conformers used in the training set can be generated by a range of methods, including informed scans and classical MD simulations. This opens the possibility that machine-learning methods, which complement all exist-

8



Figure 3. Energy errors of ML-HK along classical MD trajectories. PBE values in blue, ML-HK values in red. **a.** A 20 ps classical trajectory of benzene. **b.** A 20 ps classical trajectory of ethane.

ing approaches to functional approximation, could become a new and very different approach to this problem, with the potential to greatly reduce the computational cost of routine DFT calculations.

Our new method, directly learning the Hohenberg-Kohn density-potential map, overcomes a key bottleneck in previous methodologies that arises in 3-D. Our approach avoids solving an intermediate more general problem (the gradient descent) to find the solution of the more specific problem (finding the ground-state density). This is called transductive inference by the machine learning community and is thought to be key to successful statistical inference methods [45]. Following a direct prediction approach with the ML-HK map increases the accuracy consistently on both 1-D examples and 3-D molecules. We are also able to learn density models that outperform energy models trained on much more data. This quantitative observation allows us to conclude that learning density models is much easier than learning energy models. Such a finding should be no surprise to practitioners of the art of functional construction (see, e.g., [27]), but the present work quantifies this observation using standard statistical methods. As the ML-HK map accurately reflects the training densities, more exact methods could also be used to generate the training set densities for functional development.

We have also derived a way to use basis functions to make the approach computationally feasible. This makes it easier to integrate the method into existing DFT codes. Another advantage is the possibility to take the innate structure of the densities into account, i.e. spatial correlations are preserved by using low frequency basis functions. Again, this fits with the intuition of experienced practitioners in this field, but here we have quantified this in terms of machine-learned functionals.

Direct prediction of energies (e.g., the ML-KS map) always has the potential to lead to conceptually easier methods. But such methods must also abandon the insights and effects that have made DFT a practical and usefully accurate tool over the past half century. Many usefully accurate DFT approximations already exist, and the corrections to such approximations can be machinelearned in precisely the same way as the entire functional has been approximated here[23]. If machine-learning corrections require less data, the method becomes more powerful by taking advantage of existing successes. Furthermore, existing theorems, such as the viral theorem [46], might also be used to directly construct the kinetic energy functional from an ML-HK map. In the case of orbital-dependent functionals, such as meta-GGA's or global hybrids, the method presented here must be extended to learn, e.g., the full density matrix instead of just the density.

We also note that, for all the 3-D calculations shown here, we machine-learned E[n], the entire energy (not just the kinetic energy), which includes some densityfunctional approximation for XC. But, with a quantum chemical code, we could have trained on much more ac-



Figure 4. The precision of our density predictions using the Fourier basis for ML-HK for the molecular plane of benzene. The plots show **a**. the difference between the valence density of benzene when using PBE and LDA functionals at the PBE optimized geometry. **b**. error introduced by using the Fourier basis representation. **c**. error introduced by the $n^{\text{ML}}[v]$ density fitting (a.-c. on same color scale). **d**. the total PBE valence density **e**. the density differences along a 1-D cut of a.-c. **f**. the density error introduced with the ML-HK map (same data, but different scale, as in c.).

curate quantum chemical densities and energies. Thus, the ML-HK maps in principle allow the construction of (nearly) exact density functionals for molecular systems, with the potential to significantly reduce the computational cost of quantum chemistry based MD simulations. All this provides useful directions in which to expand on the results shown here.

METHODS

Kohn-Sham Density Functional Theory (KS-DFT) is a computational electronic structure method that determines the properties of many-body systems by using functionals of the electron density. The foundation is the Hohenberg-Kohn theorem[47] that establishes a one-to-one relationship between potential and density, i.e. at most one potential can give rise to a given groundstate density.

Kohn-Sham DFT avoids direct approximation of many body effects by imagining a fictitious system of noninteracting electrons with the same density as the real one[1]. Its accuracy is limited by the accuracy of existing approximations to the unknown exchange-correlation energy, while its computational bottleneck is the solution of the Kohn-Sham equations that describe the noninteracting particles.

Here, 3-D DFT calculations for ML models are performed with the Quantum ESPRESSO code[38] using the PBE exchange-correlation functional[35] and projector augmented waves (PAWs)[36, 37] with Troullier-Martin pseudization for describing the ionic cores[48]. All molecules are simulated in a cubic box (L = 20 bohr) with a wave function cutoff of 90 Ry. The 1-D dataset is taken from Snyder *et al.* [20].

Kernel Ridge Regression (KRR)[49, 50] is a machine learning method for regression. It is a kernelized version of Ridge Regression which minimizes the least squares error and applies an ℓ_2 (Tikhonov) regularization. Let $x_1, \ldots, x_m \in \mathbb{R}^d$ be the training data points and let $\mathbf{Y} = (y_1, \ldots, y_m)^T$ be their respective labels. KRR then optimizes



Figure 5. **a.** Energy errors of ML-HK along a 0.25 ps *ab initio* MD trajectory of malonaldehyde. PBE values in blue, ML-HK values in red. The ML model correctly predicts energies during a proton transfer in frames 7 to 15 without explicitly including these geometries in the training set. **b.** Energy errors of ML-HK along a 1 ps MD trajectory of malonaldehyde *generated by the ML-HK model*. ML-HK values in red, PBE values of trajectory snapshots in blue.

$$\min_{\alpha} \sum_{i=1}^{m} \left| y_i - \sum_{j=1}^{m} \alpha_j k(x_i, x_j) \right|^2 + \lambda \boldsymbol{\alpha}^{\mathsf{T}} \mathbf{K} \boldsymbol{\alpha}$$
(13)

where k is the kernel function and λ is a regularization parameter. **K** is the kernel matrix with $\mathbf{K}_{ij} = k(x_i, x_j)$. It admits an analytical solution

$$\boldsymbol{\alpha} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{Y}.$$
 (14)

Most popular is the Gaussian (radial basis function) kernel which allows to find a smooth non-linear model function in input space that corresponds to a linear function in an infinite dimensional feature space[28].

For the ML-HK map, the canonical error is given by the \mathcal{L}_2 distance between predicted and true densities

$$e(\boldsymbol{\beta}) = \sum_{i=1}^{M} \|n_i - n^{\mathrm{ML}}[v_i]\|_{\mathcal{L}_2}$$
(15)

$$=\sum_{i=1}^{M} \left\| n_{i} - \sum_{l=1}^{L} \sum_{j=1}^{M} \beta_{j}^{(l)} k(v_{i}, v_{j}) \phi_{l} \right\|_{\mathcal{L}_{2}}.$$
 (16)

The ML model coefficients $\beta^{(l)}$ can be optimized independently for each basis coefficient l via

$$\boldsymbol{\beta}^{(l)} = \left(\mathbf{K}_{\sigma^{(l)}} + \lambda^{(l)} \mathbf{I} \right)^{-1} \mathbf{u}^{(l)}, \quad l = 1, \dots, L.$$
(17)

Cross-validation. Note that all model parameters and hyper-parameters are estimated on the training set; the hyper-parameter choice makes use of standard crossvalidation procedures (see Hansen *et al.* [12]). Once the model is fixed after training, it is applied unchanged outof-sample.

Exact calculations. Relative energy errors of the ML models trained on KS-DFT calculations are determined by comparing to accurate energies from the Molpro Quantum Chemistry Software[51] using the Full Configuration Interaction method for H_2 and CCSD(T)[52] for H_2O .

Molecular Dynamics (MD). For benzene, ethane, and malonaldehyde, GAFF parameters[41] were assigned using the AmberTools package[53]. Geometry optimizations were performed using MP2/6-31g(d) in Gaussian09[54]. Atomic charge assignments are from RESP fits to HF/6-31g(d) calculations at optimized geometries[55] and two additional rotational conformers for ethane.

For the three larger molecules, classical isothermal MD simulations were run using the PINY_MD package[42] with massive Nosé-Hoover chain (NHC) thermostats[56]



Figure 6. The extent of the malonaldehyde conformers generated by all MD methods. a) The training set of 2,000 representative conformers selected from the classical MD trajectories (red points) by K-means sampling. Test points from an ab initio MD trajectory (green) and the independently generated MD trajectory using the ML-HK model (blue) sample the same coordinate space (offset from the molecular plane for clarity). b) A closer view of the region outlined with a dashed box for the ab initio (green) and ML-HK (blue) trajectories.

for atomic degrees of freedom (length = 4, τ = 20 fs, Suzuki-Yoshida order = 7, multiple time step = 4) and a time step of 1 fs. The r-RESPA multiple time step approach[57] was employed to compute rapidly varying forces more frequently (torsions every 0.5 fs, bonds/bends every 0.1 fs). Systems were equilibrated for 100 ps before collecting snapshots every 100 fs from 1 ns trajectories. Snapshots were aligned to a reference molecule prior to DFT calculations for the ML model. For malonaldehyde, the ML training set geometries were selected from trajectories for both enol tautomers as the GAFF force field does not permit changes in chemical bond environments.

For malonaldehyde, an additional Born-Oppenheimer MD simulation using DFT was run using the QUICK-STEP package[58] in CP2K v. 2.6.2[59]. The PBE exchange-correlation functional[35] was used in the Gaussian and plane wave (GPW) scheme[60] with DZVP-MOLOPT-GTH (m-DZVP) basis sets[61] paired with the appropriate dual-space GTH pseudopotentials[62] optimized for the PBE functional[63]. Wave functions were converged to 1E-7 Hartree using the orbital transformation method[64] on a multiple grid (n = 5) with a cutoff of 900 Ry for the system in a cubic box (L = 20 bohr). A temperature of 300 K was maintained using massive NHC thermostats[56] (length = 4, $\tau = 10$ fs, Suzuki-Yoshida order = 7, multiple time step = 4) and a time step of 0.5 fs.

In order to generate the MD trajectory with the ML-HK model, we used the Atomistic Simulation Environment (ASE) [65] with a 0.5 fs timestep and a temperature of 300 K maintained via a Langevin thermostat. A thermostat friction value of 0.01 atomic units (0.413 fs⁻¹) was chosen to reproduce the fluctuations in C atoms observed for the DFT-based trajectory (see ESI). In this proof-of-concept work, atomic forces were calculated using central finite differences, with $\epsilon = 0.001$ Å chosen to conserve the total energy during the simulation. The last 1 ps of a 4 ps trajectory was used to evaluate the performance of the ML-HK model.

DATA AVAILABILITY

All datasets used in this work are available at http: //quantum-machine.org/datasets/.

ACKNOWLEDGMENTS

We thank J.C. Snyder for early discussions and H. Glawe for helpful guidance regarding the 3-D reference computations. We thank IPAM at UCLA for repeated hospitality. Work at UC Irvine supported by NSF CHE-1464795. KRM and FB thank the Einstein Foundation for generously funding the ETERNAL project. This work was supported by Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (No. 2017-0-00451). Work at NYU supported by the U.S. Army Research Office under contract/grant number W911NF-13-1-0387 (MET and LV). Ab initio trajectory was run using High Performance Computing resources at NYU. Other DFT simulations were run using High Performance Computing resources at MPI Halle.

AUTHOR CONTRIBUTIONS

FB performed DFT calculations and ML experiments. LV performed classical and ab initio MD simulations. LL performed FCI and CCSD(T) calculations. KB and KRM initiated the work and contributed to the theory and experiments. All authors contributed to the manuscript. * to whom correspondence should be addressed.

- [1] W. Kohn and L. J. Sham, Phys. Rev. 140, A1133 (1965).
- [2] A. Pribram-Jones, D. A. Gross, and K. Burke, Annual Review of Physical Chemistry 66, 283 (2015).
- [3] J. P. Perdew, Phys. Rev. B 33, 8822 (1986).
- [4] A. D. Becke, The Journal of Chemical Physics 98, 5648 (1993).
- [5] "Materials genome initiative for global competitiveness," (2011).
- [6] A. Jain, G. Hautier, C. J. Moore, S. P. Ong, C. C. Fischer, T. Mueller, K. A. Persson, and G. Ceder, Computational Materials Science 50, 2295 (2011).
- [7] Z. D. Pozun, K. Hansen, D. Sheppard, M. Rupp, K.-R. Müller, and G. Henkelman, The Journal of Chemical Physics 136, 174101 (2012).
- [8] R. T. McGibbon and V. S. Pande, Journal of Chemical Theory and Computation 9, 2900 (2013).
- [9] T. L. Fletcher, S. J. Davie, and P. L. Popelier, Journal of chemical theory and computation **10**, 3708 (2014).
- [10] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, Phys. Rev. Lett. **108**, 058301 (2012).
- [11] G. Hautier, C. C. Fischer, A. Jain, T. Mueller, and G. Ceder, Chem. Mater. 22, 3762 (2010).
- [12] K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, J. Chem. Theory Comput. 9, 3404 (2013).
- [13] K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K.-R. Müller, and E. K. U. Gross, Phys. Rev. B 89, 205118 (2014).
- [14] K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, The Journal of Physical Chemistry Letters 6, 2326 (2015), pMID: 26113956.
- [15] K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, Nat. Commun. 8, 13890 (2017).
- [16] J. Behler and M. Parrinello, Phys. Rev. Lett. 98, 146401 (2007).
- [17] A. Seko, A. Takahashi, and I. Tanaka, Phys. Rev. B 90, 024101 (2014).
- [18] Z. Li, J. R. Kermode, and A. De Vita, Phys. Rev. Lett. 114, 096405 (2015).
- [19] S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, Sci. Adv. 3, e1603015 (2017).
- [20] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, Phys. Rev. Lett. 108, 253002 (2012).
- [21] J. C. Snyder, M. Rupp, K. Hansen, L. Blooston, K.-R. Müller, and K. Burke, J. Chem. Phys **139**, 224104 (2013).
- [22] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, International Journal of Quantum Chemistry 116, 819 (2016).
- [23] L. Li, T. E. Baker, S. R. White, and K. Burke, Phys. Rev. B 94, 245129 (2016).
- [24] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello,

L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch, Journal of Physics: Condensed Matter **21**, 395502 (19pp) (2009).

- [25] J. C. Snyder, S. Mika, K. Burke, and K.-R. Müller, in *Empirical Inference*, edited by B. Schölkopf, Z. Luo, and V. Vovk (Springer Berlin Heidelberg, 2013) pp. 245–259.
- [26] J. C. Snyder, M. Rupp, K.-R. Müller, and K. Burke, Int. J. Quantum Chem. **115**, 1102 (2015).
- [27] R. F. Ribeiro, D. Lee, A. Cangi, P. Elliott, and K. Burke, Phys. Rev. Lett. **114**, 050401 (2015).
- [28] K.-R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf, IEEE Trans. Neural Netw. 12, 181 (2001).
- [29] M.-C. Kim, E. Sim, and K. Burke, Phys. Rev. Lett. 111, 073003 (2013).
- [30] M.-C. Kim, E. Sim, and K. Burke, The Journal of Chemical Physics 140, 18A528 (2014).
- [31] M.-C. Kim, H. Park, S. Son, E. Sim, and K. Burke, J. Phys. Chem. Lett. 6, 3802 (2015).
- [32] R. M. Dreizler and E. K. U. Gross, Density Functional Theory: An Approach to the Quantum Many-Body Problem (Springer-Verlag Berlin Heidelberg, 1990).
- [33] B. Schölkopf, S. Mika, C. Burges, P. Knirsch, K.-R. Müller, G. Rätsch, and A. Smola, IEEE Trans. Neural Netw. 10, 1000 (1999).
- [34] B. Schölkopf, A. Smola, and K.-R. Müller, Neural Comput. 10, 1299 (1998).
- [35] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. 77, 3865 (1996).
- [36] G. Kresse and D. Joubert, Phys. Rev. B 59, 1758 (1999).
- [37] P. E. Blöchl, Phys. Rev. B 50, 17953 (1994).
- [38] 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. d. Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch, J. Phys.: Condens. Matter **21**, 395502 (2009).
- [39] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, Phys. Rev. Lett. **104**, 136403 (2010).
- [40] M. J. D. Powell, The Computer Journal 7, 155 (1964).
- [41] J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, Journal of Computational Chemistry 25, 1157 (2004).
- [42] M. E. Tuckerman, D. Yarne, S. O. Samuelson, A. L. Hughes, and G. J. Martyna, Computer Physics Communications 128, 333 (2000).
- [43] J. P. Perdew and A. Zunger, Phys. Rev. B 23, 5048 (1981).
- [44] M. E. Tuckerman and D. Marx, Phys. Rev. Lett. 86, 4946 (2001).
- [45] V. Vapnik, The Nature of Statistical Learning Theory, Information Science and Statistics (Springer, 2000).
- [46] A. Cangi, D. Lee, P. Elliott, K. Burke, and E. K. U. Gross, Phys. Rev. Lett. **106**, 236404 (2011).
- [47] P. Hohenberg and W. Kohn, Phys. Rev. 136, B864 (1964).
- [48] N. Troullier and J. L. Martins, Phys. Rev. B 43, 1993 (1991).
- [49] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning Data Mining, Inference, and Pre*diction, 2nd ed., Springer Series in Statistics (Springer,

2009).

- [50] K. Vu, J. C. Snyder, L. Li, M. Rupp, B. F. Chen, T. Khelif, K.-R. Müller, and K. Burke, International Journal of Quantum Chemistry 115, 1115 (2015).
- [51] H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, M. Schütz, *et al.*, "Molpro, version 2015.1, a package of ab initio programs," (2015).
- [52] T. B. Adler, G. Knizia, and H.-J. Werner, Journal of Chemical Physics 127, 221106 (2007).
- [53] J. Wang, W. Wang, P. A. Kollman, and D. A. Case, Journal of the American Chemical Society 222, U403 (2001).
- [54] M. Frisch, G. Trucks, H. B. Schlegel, G. Scuseria, M. Robb, J. Cheeseman, G. Scalmani, V. Barone, B. Mennucci, G. Petersson, et al., "Gaussian 09," (2009).
- [55] C. I. Bayly, P. Cieplak, W. Cornell, and P. A. Kollman, The Journal of Physical Chemistry 97, 10269 (1993).
- [56] G. J. Martyna, M. L. Klein, and M. Tuckerman, The Journal of Chemical Physics 97, 2635 (1992).

- [57] M. Tuckerman, B. J. Berne, and G. J. Martyna, The Journal of Chemical Physics 97, 1990 (1992).
- [58] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, Computer Physics Communications 167, 103 (2005).
- [59] J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, Wiley Interdisciplinary Reviews: Computational Molecular Science 4, 15 (2014).
- [60] G. Lippert, J. Hutter, and M. Parrinello, Molecular Physics 92, 477 (2010).
- [61] J. VandeVondele and J. Hutter, The Journal of Chemical Physics 127, 114105 (2007).
- [62] S. Goedecker, M. Teter, and J. Hutter, Phys. Rev. B 54, 1703 (1996).
- [63] M. Krack, Theoretica Chimica Acta 114, 145 (2005).
- [64] J. VandeVondele and J. Hutter, The Journal of Chemical Physics 118, 4365 (2003).
- [65] S. R. Bahn and K. W. Jacobsen, Comput. Sci. Eng. 4, 56 (2002).