

**Deep learning and density-functional theory**

Kevin Ryczko\*

*Department of Physics, University of Ottawa, Ottawa, Ontario K1N 6N5, Canada*

David A. Strubbe

*Department of Physics, University of California Merced, Merced, California 95343, USA*

Isaac Tamblyn†

*National Research Council of Canada and Department of Physics, University of Ottawa, Ottawa, Ontario K1N 6N5, Canada*

(Received 28 February 2019; published 21 August 2019)

We show that deep neural networks can be integrated into, or fully replace, the Kohn-Sham density functional theory (DFT) scheme for multielectron systems in simple harmonic oscillator and random external potentials with no feature engineering. We first show that self-consistent charge densities calculated with different exchange-correlation functionals can be used as input to an extensive deep neural network to make predictions for correlation, exchange, external, kinetic, and total energies simultaneously. Additionally, we show that one can make all of the same predictions with the external potential rather than the self-consistent charge density, which allows one to circumvent the Kohn-Sham DFT scheme altogether. We then show that a self-consistent charge density found from a nonlocal exchange-correlation functional can be used to make energy predictions for a semilocal exchange-correlation functional. Lastly, we use a deep convolutional inverse graphics network to predict the charge density given an external potential for different exchange-correlation functionals and assess the viability of the predicted charge densities. This work shows that extensive deep neural networks are generalizable and transferable given the variability of the potentials (maximum total energy range  $\approx 100$  Ha) because they require no feature engineering and because they can scale to an arbitrary system size with an  $O(N)$  computational cost.

DOI: [10.1103/PhysRevA.100.022512](https://doi.org/10.1103/PhysRevA.100.022512)**I. INTRODUCTION**

Kohn-Sham (KS) density functional theory (DFT) [1] is the standard theoretical tool to study nanoscale systems. Despite its success, DFT calculations for atomistic systems containing tens of thousands to millions of atoms are exceptionally demanding from a computational perspective and are rare in the literature. Machine-learning techniques can replace conventional DFT calculations to overcome this computational barrier. Machine-learning models are ideal because they rival the accuracy of the method they are trained on, but can be less demanding to evaluate from a computational standpoint. There have been many reports where artificial neural networks (ANNs) have been used to represent potential-energy surfaces to accelerate electronic structure calculations [2–9]. These reports focus on feature engineering or defining some abstract representation of atomistic systems allowing one to use an ANN. Instead, we focus our review on reports that avoid feature engineering and utilize the electron density in conjunction with machine learning. More specifically, machine learning has become a popular choice to represent energy functionals in DFT [10–14] or to completely circumvent the KS scheme [15,16]. In deep learning, the machine-learning model learns the hierarchical features during training rather than inputting

abstract representations. Due to the large number of tunable parameters in deep neural networks (DNNs) that may include a variety of layers (i.e., convolutional, fully connected, max pooling, etc.), there must be thousands (if not hundreds of thousands) of training examples to find a stable minima with an acceptable accuracy. Generating these training examples is a computationally expensive task, but a trained DNN can evaluate a given quantity at a fraction of the cost compared to the original method.

An alternative approach that has been taken recently by Brockherde *et al.* [15] is to focus more on uniformly sampling the space that a machine-learning model will eventually predict and to use traditional machine learning with far fewer tunable parameters. This approach was successful in predicting KS-DFT total energies and charge densities in one dimension (1D) for random Gaussian potentials and for small molecules in three dimensions (3D). Due to the use of kernel ridge regression (KRR), they were able to achieve an acceptable accuracy with a relatively small number of training examples. Unfortunately, KRR is known to have poor scaling with respect to the number of training examples, making it difficult to train with a large (and more diverse) set of training examples.

In KS-DFT, one of the contributions to the total energy is the noninteracting kinetic energy. Before the KS scheme was realized, Hohenberg and Kohn [17] postulated the formalism for an interacting kinetic-energy functional of the density. An analytic expression for the exact interacting kinetic-energy

\*kryczko@uottawa.ca

†isaac.tamblyn@nrc.ca

functional with respect to the electron density is unknown. This is one of the major downfalls of orbital-free (OF) DFT, where all energy contributions are explicitly written in terms of the electron density. This shortcoming provides motivation to construct an approximate functional of the density with a machine-learning model. In a report done by Yao *et al.* [18], a convolutional neural network (CNN) was used to represent the kinetic-energy functional in the OF-DFT total-energy expression for various hydrocarbons. Their data-generation process consisted of performing KS-DFT and collecting the charge density along with the KS noninteracting kinetic energy. The charge density was then used as input to the CNN with the KS noninteracting kinetic energy as the label. With this representation, they were able to successfully reproduce potential-energy surfaces when compared to the true KS potential-energy surfaces. In another report by Snyder *et al.* [13], they were able to use a machine-learning model to make kinetic-energy predictions given a charge density for a diatomic molecule. Using their framework, they were able to accurately dissociate the diatomic molecule and compute forces suggesting that *ab initio* molecular dynamics could eventually be done via machine-learning methods.

When representing the kinetic energy with a machine-learning model in the OF scheme, one then becomes concerned with calculating the functional derivative of the machine-learning model with respect to the density. In a report from Li *et al.* [12], they showed that there is a trade-off between accuracy and numerical noise when taking the functional derivative of a machine-learning model. Brockherde *et al.* [15] avoided this issue by training a machine-learning model to learn the mapping between the potential and the electron density, avoiding the functional derivative.

In another recent report by Kolb *et al.* [14], a software package was developed to combine artificial neural networks with electronic structure calculations and molecular dynamics engines. Using their software, they were able to show that artificial neural networks can be used to make predictions with the electronic charge density as input and various energies as output. Specifically, they were able to predict energies and band gaps calculated at a higher level of theory from charge densities calculated at a lower level of theory. This approach is very advantageous as high-level theory calculations (i.e.,  $G_0W_0$  [19]) become quite computationally expensive for larger systems.

Although significant progress has been made incorporating machine learning and deep learning to a variety of electronic structure problems, most do not have the ability to properly handle extensive properties. In some of our past work, we introduced extensive DNNs (EDNNs) [20] to intrinsically learn extensive properties. This means that when the DNN learns the fundamental screening length scale, it can then easily scale up to massive systems in a trivially parallel manner. EDNNs work by first dividing up an image into fragments, which are called focus regions. These fragments are then padded with context regions. The context regions may also respect periodic boundary conditions. Each of these fragments can then be simultaneously passed into machine-learning models that share weights. It should be noted here that any machine-learning method that uses back propagation to minimize the loss function can be used. Finally, the outputs

of the machine-learning models are then summed, yielding the final prediction from the EDNN.

In this report, we show that EDNNs have the capability to learn energy and charge density mappings that could replace some, if not all, calculations in the KS-DFT scheme. We push the frontier of what EDNNs can learn from charge densities and external potentials by calculating the self-consistent charge densities in external potentials with extreme variabilities. In previous reports that focused on small molecules [5,14,15,18], the self-consistent charge densities generated from molecular dynamics are similar and have small energy ranges (i.e.,  $\approx 31.8$  mHa [15]). We avoid small molecules (where the charge density would be localized in space) and truly challenge the ability of EDNNs to make accurate predictions across a variety of electronic environments. Quantitatively speaking, the energy range of our 10 electron calculations with our random (RND) external potentials is  $\approx 100$  Ha. This report is outlined as follows: In Sec. II, we describe our data-generation process, as well as the DNN topologies and hyperparameter selections. In Sec. III A, we show that DNNs have the capability to act as density functionals and can accurately predict the exchange, correlation, external, kinetic, and total energies *simultaneously* (Sec. III A 1). We also show that EDNNs can circumvent the KS scheme (Sec. III A 2) by mapping the external potential to all of the aforementioned energies simultaneously. Additionally, we show that EDNNs can be used in a somewhat “perturbative” manner, where we predict energies computed with semilocal or nonlocal exchange-correlation functionals from nonlocal electron densities. In Sec. III B, we show that deep convolutional inverse graphics networks (DCIGNs) can also map the external potential to the electron density, and assess the viability of the predicted electron density. Lastly, in Sec. IV, we summarize our results and consider future work that could be done with our framework. The outline of this manuscript can be seen graphically in Fig. 1.

## II. METHODS

We investigate two-dimensional (2D) electron gases within the KS-DFT framework [1] for two external potentials: simple harmonic oscillator (SHO) and RND. These external potentials have been used in a previous study [22] for direct-diagonalization, one-electron calculations. In the KS-DFT framework, one minimizes the total-energy functional,

$$E[\rho] = T[\rho] + E_{\text{ext}}[\rho] + E_{\text{Hartree}}[\rho] + E_{\text{XC}}[\rho]. \quad (1)$$

In Eq. (1),  $T$  is the noninteracting kinetic energy,  $E_{\text{ext}}$  is the energy due to the interaction of the electrons with the external potential,  $E_{\text{Hartree}}$  is the electrostatic energy describing the electron-electron interactions, and  $E_{\text{XC}}$  is the exchange-correlation energy.

Using EDNNs, we investigate the feasibility of learning the total energy as well as the individual contributions to the total energy. We therefore have trained models to predict the total, noninteracting kinetic, external, exchange, and correlation energies. The external potentials chosen for this report, as mentioned previously, are SHO and RND potentials in 2D.

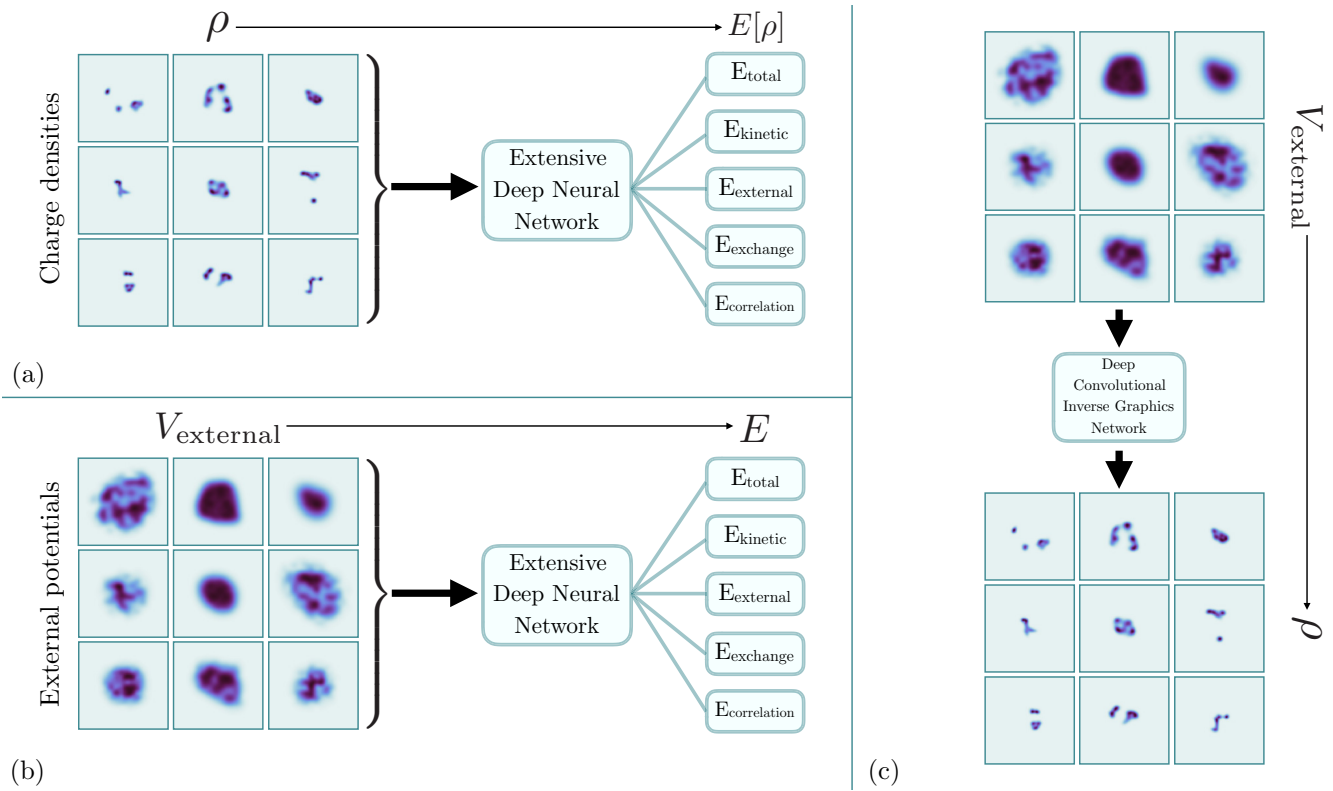


FIG. 1. A graphical representation that outlines the objectives of this report. We show that both (a) charge densities and (b) external potentials can be used as input to extensive deep neural networks (EDNNs [20]) to predict the total, kinetic, external, exchange, and correlation energies. The images shown are some of the random (RND) potentials along with the self-consistent charge density for that potential. (c) We show that deep convolutional inverse graphics networks (DCIGNs [21]) can be used to map external potentials to their respective self-consistent charge densities.

The SHO potentials take the form

$$V_{\text{ext}}(\{x_i\}) = \frac{1}{2} \sum_i^D k_i (x_i - x_{0i})^2, \quad (2)$$

where  $D$  is the dimension,  $k_i = m\omega_i^2$  is the spring constant, and  $x_{0i}$  is the shift of the potential in a given coordinate. For the RND potentials, we follow the work of Mills *et al.* [22] when generating the potentials on a grid. We refer the reader to the original manuscript [22] for more information on the

RND potential generation. Examples of the RND external potentials along with the respective charge densities can be seen in Fig. 2. To briefly summarize the process, the first step consists of generating a random binary image of 0's and 1's and then applying a Gaussian filter. The second step consists of generating a mask that is constructed with a random convex hull and an additional Gaussian blur. The mask is then applied onto the image, yielding the final result. The larger energy scale of the RND external potentials can be attributed to the length scales of the RND external potentials. The average

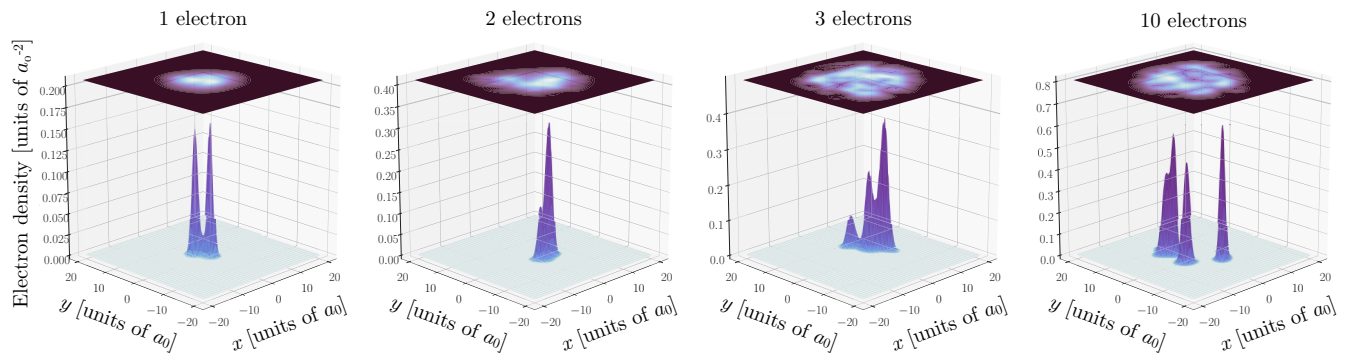


FIG. 2. Computed charge densities (3D surfaces) and random external potential-energy surfaces (2D surfaces) for typical configurations of systems with 1, 2, 3, and 10 electrons.

Gaussian kernel sizes in the external potential generation is  $\sim 3$  Bohr, whereas the average length scale of the SHO external potentials is  $\sim 12$  Bohr. Assuming that the energy scales as  $E \sim 1/r^2$ , the energy scale of the RND external potentials is 16 times larger than the SHO external potential-energy scale on average. To create datasets large enough to use DNNs, we chose to randomly sample  $k_i$  and  $x_0$ , such that  $0.01 \leq k_i \leq 0.16 \text{ Ha}/a_0^2$  (Hartree per Bohr<sup>2</sup>) and  $-8.0 \leq x_0 \leq 8.0 a_0$ . With a given selection of these variables, the external potential was then evaluated on a  $40 \times 40 a_0$  space with a  $256 \times 256$  grid point mesh.

We then chose to place either  $N = 1, 2, 3$ , or 10 electrons in the 2D space. For each choice of the number of electrons, we generated an external potential, and then performed three DFT calculations, each with different exchange-correlation functionals. We used the local density approximation (LDA) Slater exchange functional in 2D [23,24], the Perdew-Burke-Ernzerhof (PBE) generalized-gradient approximation (GGA) exchange functional in 2D [25,26] and the metageneralized-gradient approximation (MGGA) exchange functional from Pittalis *et al.* [27], in each case paired with the correlation functional of Attaccalite *et al.* [28], all provided by the LIBXC library [29,30]. Here, we do not take the orbitals or charge densities from the LDA calculations to calculate energies at the PBE or MGGA level. All of the energies for each functional are calculated independently.

All of the calculations were carried out in real space with the OCTOPUS code [31,32]. For testing, we set aside 10% of each dataset. This made for 90 000 training configurations and 10 000 testing configurations for each case of potential, number of electrons, and exchange-correlation functional. Note that the number of electrons and the type of external potential uniquely defines a dataset. Therefore, the 10% set aside for testing includes 10 000 external potentials and 30 000 charge densities (10 000 for each exchange-correlation functional choice). In addition, the labels were normalized independently such that each of the components had a range from zero to one. The calculations are summarized in Table I of Appendix A. All data used in this report are available online [37] along with the code [33] to allow for future development of featureless deep-learning-based functionals.

When constructing the EDNNs, we used a mixture of TENSORFLOW [34] and TFLearn [35] in PYTHON. For the networks topologies, we build on our previous reports [16,20] and use EDNNs where each tile of the EDNN has the same in-tile CNN used previously for predicting KS-DFT total energies of 2D hexagonal sheets [16]. For clarity, the in-tile CNN consisted of two reducing convolutional layers with kernel sizes of 3, six nonreducing convolutional layers with kernel sizes of 4, one reducing convolutional layer with a kernel size of 3, four nonreducing convolutional layers with kernel sizes of 3, a fully connected layer with 1024 neurons, and a final fully connected layer with one neuron. All of the activations used were rectified linear units. We emphasize that in our approach, we do not do any sort of feature engineering, such as past reports that use ANNs [2–7,14]. The convolutional layers in the EDNNs identify relevant features during the training process. When utilizing an EDNN, one must declare the focus and context regions which are used to “tile” up the image into fragments. To find the ideal focus

and context regions, we started by training the EDNNs on the 2D charge density to total-energy mapping as well as the 2D external potential to total-energy mapping for the 1-, 2-, 3-, and 10-electron systems for calculations done with the LDA exchange-correlation functional and the SHO external potential. We chose a variety of focus and context sizes, and found that the optimal focus and context sizes are 128 pixels for the focus size and 32 pixels for the context size. Our decision was based on a balance between accuracy and computation time. A larger focus size lowers the computation time, and a larger context size yields larger images, resulting in more neurons in the EDNNs thereby improving the accuracy of the model. For a focus of 128 pixels, we found that the accuracies were very similar for various context sizes and the choice of 32 pixels was almost arbitrary. This hyperparameter search was the most computationally demanding task for this work due to the number of models that had to be trained. While training, we used a learning rate of  $10^{-4}$  for 500 epochs when using the charge densities as input and a learning rate of  $10^{-5}$  for 500 epochs when using the external potentials as input. In both cases, we further reduced the learning rates by a factor of 10 and trained for an additional 100 epochs. For clarity, an epoch is defined to be when the weights of the network have been updated for the entire training dataset once.

### III. RESULTS

#### A. Energy predictions

##### 1. EDNNs as a functional

First, we show that EDNNs can be used as an energy functional for correlation, exchange, external, kinetic, and total energies. For the LDA, PBE, and MGGA functionals discussed in Sec. II, we used the computed self-consistent charge densities as input to an EDNN and were able to successfully predict the correlation, exchange, external, and total energies simultaneously for both SHO and RND external potentials. Starting with the models where the SHO external potentials were used in the DFT calculations, we found that the mean absolute errors for each particular case are less than 1.5 mHa. These can be seen in Table II of Appendix B. In Fig. 3, we show predicted minus true versus true for the 1- and 10-electron models with the different exchange-correlation functionals when the RND external potentials were used in the DFT calculations. In this figure, it is clear that the error of the models increases with the number of electrons. This increase in error is expected due to the increase in the range of energies and can be physically attributed to the increase of interactions in the system. Looking to Table III of Appendix C, we also observe that the mean absolute errors become larger as the complexity of the exchange-correlation functional increases. In addition to these trends, we also notice that the energy with the largest mean absolute error comes from the external-energy functional. This again can be attributed to the ranges of the various energies. The external energy has the largest range of all the energies being predicted. We also address the generalizability of the models by testing the model that was trained on 10-electron charge densities calculated with the RND external potentials and the LDA functional with 10-electron charge densities calculated with the SHO external

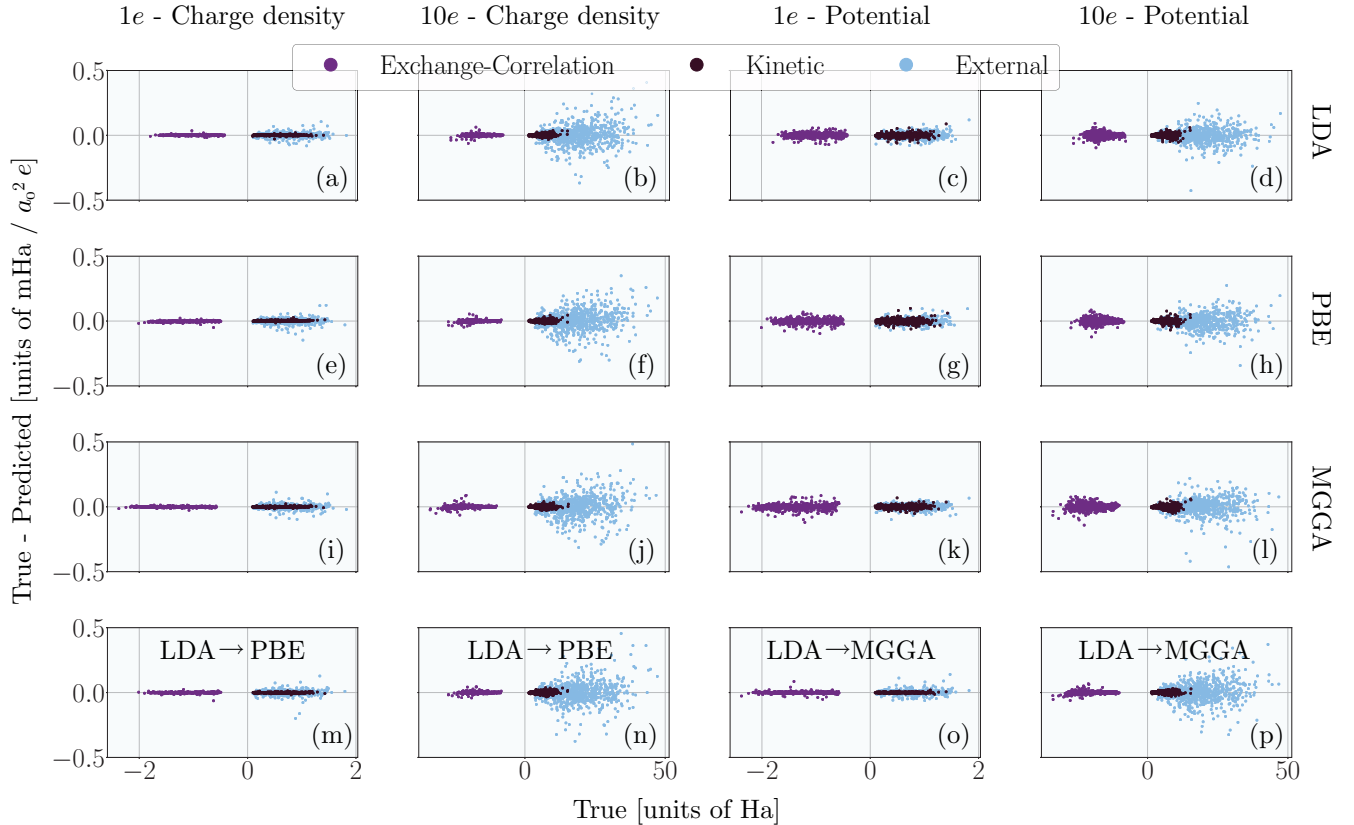


FIG. 3. True minus predicted (in  $\text{mHa}/a_0^2/\text{electron}$ ) vs true (Ha) for various models with the RND potentials. Models trained (a)–(d) with the LDA exchange-correlation functional, (e),(f) with the PBE exchange-correlation functional, and (i)–(l) with the MGGA exchange-correlation functional. (a),(e),(i) The first column is for 1-electron models where the charge densities were used as input. (b),(f),(j) The second column is for 10-electron models where the charge densities were used as input. (c),(g),(k) The third column is for 1-electron models where the external potentials were used as input. (d),(h),(l) The fourth column is for 10-electron models where the external potentials were used as input. (m)–(p) The bottom row is for models where LDA charge densities were used as input, and the labels were either (m),(n) PBE energies or (o),(p) MGGA energies. (m),(o) The 1-electron systems; (n),(p) the 10-electron systems. It should be noted that one model is predicting the correlation, exchange, external, kinetic, and total energies. We have combined the exchange-correlation error and omitted the total-energy error for clarity.

potentials and the LDA functional (and vice versa). We found in both cases that the errors increased by several orders of magnitude. This is not surprising given the different energy ranges of the datasets. On the contrary, when examining the true versus predicted plot for the model trained on the RND dataset but tested using the SHO dataset, we found that a constant shift could simply be added to substantially decrease the error. We expect that this constant shift could be easily rectified if SHO training examples are included in the training process.

DFT is a more popular choice for larger systems relative to wave-function-based methods because the exchange-correlation functionals used are computationally inexpensive relative to methods that employ exact exchange, for example. In light of this, we have trained EDNNs to predict energies at the PBE and MGGA level given a self-consistent charge density computed with the LDA exchange-correlation functional. In Fig. 3, we consider 1- and 10-electron models trained on the mapping between LDA charge densities and either PBE or MGGA energies. Similar to the results mentioned above, the mean absolute errors increase both with the number of electrons and the complexity of the exchange-correlation

functional. In Table III of Appendix C, we also notice that the highest mean absolute error is for the external and total energies. This result further suggests that there is not a fundamental problem with learning the external energy, but the larger range of energies makes it more difficult for a EDNN to handle with extreme precision. In addition, since the correlation functional is the same across all of the calculations and the same testing data were used for each case of number of electrons, we can determine if the networks are learning the correlation energy in a similar manner. In Table III, we can see that the correlation energies have similar magnitudes of error, indicating that similar correlation functional mappings are being learned as one should expect. The success of learning the energies of a more accurate exchange-correlation functional given a less accurate charge density shows promise for other applications. A future application could include learning  $G_0W_0$  quasiparticle energies from a DFT-computed self-consistent charge density, similar to the work that was completed by Kolb *et al.* [14].

A note should be made about Table III with respect to the magnitude of some of the mean absolute errors reported. In comparison to the report by Mills *et al.* [22], some of the

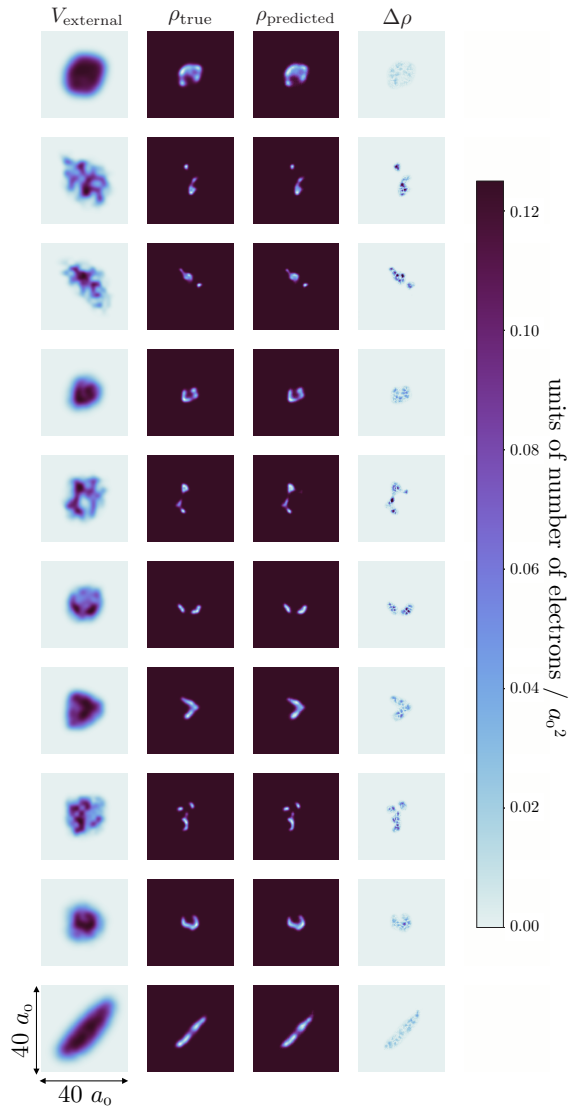


FIG. 4. Examples of the RND potentials, true charge densities ( $\rho_{\text{true}}$ ), predicted charge densities ( $\rho_{\text{predicted}}$ ), and differences between the true and predicted charge densities ( $\Delta\rho$ ). The charge densities shown here were computed with the LDA exchange-correlation functional. The color bar is for the charge density differences  $\Delta\rho$ .

mean absolute errors are larger, in some cases, by a factor of 10. In addition, the focus and context hyperparameters were optimized for the SHO external potentials. In the work of Brockherde *et al.* [15], they managed to reach chemical accuracy using three-dimensional charge densities, but the energy range of their training set was  $\approx 40$  kcal/mol (for a benzene molecule). For their best reported model with a mean absolute error of 0.28 kcal/mol, the relative mean absolute error, which we define as the mean absolute error divided by the range of the dataset, is 0.007. In our 10-electron model with the RND external potential, our energy range was  $\approx 100$  Ha (62 750.9 kcal/mol) and the mean absolute error of the total energy predictions was 78.514 mHa, yielding a relative mean absolute error of  $7.85 \times 10^{-4}$ .

## 2. Circumventing Kohn-Sham DFT

In addition to using EDNNs as a functional, it is arguably more convenient to train a EDNN to learn the mapping between the external potential and the contributing energies of that system. It is more convenient because it avoids calculating a self-consistent charge density with the KS scheme. We have trained EDNNs to predict the exchange, correlation, external, kinetic, and total energy simultaneously using the external potential as input rather than the charge density. Again, in Fig. 3, we show true minus predicted versus true for the correlation, exchange, external, kinetic, and total energies for the RND external potentials. Here, it is evident that the charge density is more optimal as an input to an EDNN for the 1-electron systems. There is much more spread in the distribution when using external potentials as input compared to charge densities. For 10 electrons, this is not the case. Looking to Table III of Appendix C, we can see that for 1, 2, and 3 electrons, no matter what choice of exchange-correlation functional, it is less difficult to learn the mapping between  $\rho \rightarrow E$  than  $V \rightarrow E$ . The mean absolute errors are lower for all energies. In the case of 10 electrons, the mean absolute errors in the external and total energies are lower for the models that have potentials as input. Although the errors are lower for the external and total energies, the mean absolute errors for correlation, exchange, and kinetic energies are larger. When training a model on a set of energies, there is a balance between the errors of the energies since the loss function depends on the sum over the mean-squared errors between the true and predicted energies. In the case of using charge densities as input to the EDNN, we found the exchange, correlation, and kinetic energies can be predicted with much better accuracy than the external or total energies. In the case of using potentials as input to the EDNN, we found that there is more of a balance of accuracy between the different energies being predicted.

## B. Image predictions with deep convolutional inverse graphics networks

In both KS-DFT and OF-DFT, the self-consistent charge density is the central quantity that one is interested in calculating. Once one has the charge density, most other quantities can be calculated in a straightforward manner. In this section, we address the viability of using DCIGNs to map the external potential to the self-consistent charge density in 2D for the RND potentials with the LDA, PBE, and MGGA exchange-correlation functionals. DCIGNs were recently introduced in the literature [21] and have a similar topology to autoencoders [36]. The DCIGN that we have used has four reducing convolutional layers, three nonreducing convolutional layers, and four deconvolutional layers such that the output image has the same dimensionality as the input image. This topology differs slightly from the original work on DCIGNs [21], where a fully connected layer would replace our three nonreducing convolutional layers. Additionally, our DCIGN is deterministic. In the original work [21], random noise is introduced in the decoder to create a nondeterministic generative model. All of our convolutional layers use a kernel size of 3 with rectified linear unit activations. We used a learning rate of  $10^{-5}$  while training for 500 epochs and dropped the learning

rate by a factor of 10 before training for an additional 100 epochs. For this discussion, we focus solely on the 10-electron calculations with the RND external potentials. We argue that these are the most challenging calculations to train with a DCIGN, and can therefore safely assume that the less complex calculations would be successful given the success of the most complex cases. In Fig. 4, we show some of the predictions that the DCIGN made for 10-electron calculations with the LDA exchange-correlation functional. There is a remarkable resemblance between the true ( $\rho_{\text{true}}$ ) and predicted ( $\rho_{\text{predicted}}$ ) charge densities. The DCIGN is capable of handling the extreme variability of the complex shapes and is capable of handling the cases where the charge density is not isolated to one region of space. From a qualitative perspective, the DCIGN makes accurate predictions of the charge densities given RND external potentials.

Normally, when addressing the viability of a machine-learning model from a quantitative perspective, one considers the mean absolute error on the test set. We argue that a more rigorous test for  $\rho_{\text{predicted}}$  would be the mean absolute error of the energies associated with  $\rho_{\text{predicted}}$ . We therefore take  $\rho_{\text{predicted}}$  and renormalize them such that  $\int d\mathbf{r} \rho(\mathbf{r})_{\text{predicted}} = 10$ . Afterwards, we use the renormalized  $\rho_{\text{predicted}}$  as input to a subset of the models described in Sec. III A 1. We then compare the energies predicted from  $\rho_{\text{predicted}}$  with the true energies. In Table III of Appendix C, we show the mean absolute errors between the true and predicted energies for the different exchange-correlation functionals. When comparing the mean absolute errors of the predicted energies for  $\rho_{\text{predicted}}$  with the energy predictions made from the  $\rho_{\text{true}}$ , the minimal difference was seen for the correlation energies with a value of  $\sim 6$  mHa. This was true for all exchange-correlation functionals considered in this work. The maximal difference between the mean absolute errors when comparing the energy predictions of  $\rho_{\text{predicted}}$  and  $\rho_{\text{true}}$  was the total energy, which was  $\sim 20$  mHa. Again, this is true for all exchange-correlation functionals considered. In addition to this metric, we also report the mean absolute error for a model mapping  $\rho_{\text{predicted}}$  to the true energies. We find an increase in the errors in comparison to the models that map the true charge densities to the true energies, but the errors are comparable to the errors when we evaluate  $\rho_{\text{predicted}}$  with the model trained on the true charge densities. Training with the charge densities

and energies as labels yields similar results. We also report the mean absolute density driven error (DDE) using a similar definition as Brockherde *et al.* [15] in Table III. We find similar trends in the DDEs when comparing them to the mean absolute errors of  $\rho_{\text{predicted}}$ . For more information on our definition of the mean absolute DDEs, see Appendix D. In addition, to compare with Brockherde *et al.* [15], our relative mean absolute error is a factor of  $\approx 7$  times smaller.

#### IV. CONCLUSION

In conclusion, we have shown that extensive deep neural networks (EDNNs) and deep convolutional inverse graphics networks (DCIGNs) can be used alongside or replace conventional Kohn-Sham density-functional-theory (KS-DFT) calculations. For both the random (RND) and simple harmonic oscillator (SHO) external potentials, EDNNs have the capability to make highly accurate energy predictions using both the charge densities and the external potentials as input for correlation, exchange, external, kinetic, and total energy simultaneously (dataset is available online [37]). In addition, we have shown that DCIGNs have the capability to predict charge densities given an external potential. Qualitatively speaking, the predicted charge densities are remarkably similar to the true charge densities. Quantitatively speaking, the relative mean absolute errors were found to be smaller than previous, state-of-the-art work [15]. The results of this report show promise for future application in two regards. First, this framework has the capability to make predictions of higher-level theory calculations given a lower-level theory charge density similar to a previous report [14]. Second, both EDNNs and DCIGNs can be used to calculate energies covering a large range of electronic environments to a high level of accuracy.

#### ACKNOWLEDGMENTS

K.R. and I.T. acknowledge funding from the Natural Sciences and Engineering Research Council of Canada, and Compute Canada and SOSCIP for computational resources. K.R. and I.T. would also like to thank NVIDIA for a faculty hardware grant. D.A.S. acknowledges University of California, Merced start-up funding.

#### APPENDIX A: MORE DETAILS ON THE GENERATED DATA

TABLE I. Summary of the calculations that were used for training and testing the deep-learning models.  $N$  is the number of electrons,  $V_{\text{ext}}$  is the external potential chosen (see text), and  $V_X + V_C$  are the exchange-correlation potentials chosen. Note that the combination of the number of electrons and external potential produces a unique dataset. For example, the 3-electron systems in the RND potentials have 100 000 external potentials but contributes 300 000 calculations due to the use of different exchange-correlation functionals.

$N$	$V_{\text{ext}}$	$V_X + V_C$	No. calculations
1	SHO	LDA + LDA	100 000
1	SHO	PBE + LDA	100 000
1	SHO	MGGA + LDA	100 000

TABLE I. (*Continued.*)

$N$	$V_{\text{ext}}$	$V_X + V_C$	No. calculations
1	RND	LDA + LDA	100 000
1	RND	PBE + LDA	100 000
1	RND	MGGA + LDA	100 000
2	SHO	LDA + LDA	100 000
2	SHO	PBE + LDA	100 000
2	SHO	MGGA + LDA	100 000
2	RND	LDA + LDA	100 000
2	RND	PBE + LDA	100 000
2	RND	MGGA + LDA	100 000
3	SHO	LDA + LDA	100 000
3	SHO	PBE + LDA	100 000
3	SHO	MGGA + LDA	100 000
3	RND	LDA + LDA	100 000
3	RND	PBE + LDA	100 000
3	RND	MGGA + LDA	100 000
10	SHO	LDA + LDA	100 000
10	SHO	PBE + LDA	100 000
10	SHO	MGGA + LDA	100 000
10	RND	LDA + LDA	100 000
10	RND	PBE + LDA	100 000
10	RND	MGGA + LDA	100 000
Total			2 400 000

## APPENDIX B: MEAN ABSOLUTE ERRORS WITH THE SIMPLE HARMONIC OSCILLATOR EXTERNAL POTENTIALS

TABLE II. Mean absolute errors (in mHa per electron) and root-mean-squared errors (in parentheses) for models trained in this report for the SHO potentials. The abbreviations  $\rho$  and  $V_{\text{ext}}$  are charge density and potential, respectively. The arrows (i.e., LDA  $\rightarrow$  PBE) indicate that the charge density used as input to the DNN was calculated using the LDA exchange-correlation functional, but the labels (energies) were calculated using another exchange-correlation functional.

$N_{\text{electrons}}$	Input	Functional	$E_{\text{correlation}}$	$E_{\text{exchange}}$	$E_{\text{external}}$	$E_{\text{kinetic}}$	$E_{\text{total}}$
1	$\rho$	LDA	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.2 (0.2)
2	$\rho$	LDA	0.1 (0.1)	0.1 (0.2)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
3	$\rho$	LDA	0.1 (0.1)	0.1 (0.2)	0.2 (0.3)	0.1 (0.2)	0.5 (0.6)
10	$\rho$	LDA	0.1 (0.1)	0.2 (0.2)	0.3 (0.6)	0.2 (0.3)	0.9 (1.4)
1	$\rho$	PBE	0.1 (0.1)	0.2 (0.2)	0.1 (0.2)	0.1 (0.1)	0.2 (0.3)
2	$\rho$	PBE	0.1 (0.1)	0.2 (0.2)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
3	$\rho$	PBE	0.1 (0.2)	0.2 (0.3)	0.3 (0.4)	0.2 (0.2)	0.8 (0.9)
10	$\rho$	PBE	0.1 (0.1)	0.2 (0.2)	0.4 (0.6)	0.2 (0.2)	0.9 (1.5)
1	$\rho$	MGGA	0.1 (0.1)	0.2 (0.3)	0.1 (0.2)	0.1 (0.1)	0.3 (0.3)
2	$\rho$	MGGA	0.2 (0.2)	0.3 (0.4)	0.2 (0.2)	0.1 (0.2)	0.4 (0.5)
3	$\rho$	MGGA	0.1 (0.1)	0.2 (0.3)	0.2 (0.2)	0.1 (0.2)	0.4 (0.5)
10	$\rho$	MGGA	0.2 (0.3)	0.4 (0.5)	0.5 (0.8)	0.3 (0.4)	1.3 (1.9)
1	$V_{\text{ext}}$	LDA	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.2 (0.2)
2	$V_{\text{ext}}$	LDA	0.1 (0.1)	0.1 (0.2)	0.1 (0.1)	0.1 (0.1)	0.2 (0.3)
3	$V_{\text{ext}}$	LDA	0.1 (0.1)	0.1 (0.2)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
10	$V_{\text{ext}}$	LDA	0.1 (0.2)	0.2 (0.3)	0.3 (0.8)	0.2 (0.3)	0.8 (1.1)
1	$V_{\text{ext}}$	PBE	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.1 (0.1)	0.1 (0.2)
2	$V_{\text{ext}}$	PBE	0.1 (0.1)	0.1 (0.2)	0.1 (0.1)	0.0 (0.1)	0.2 (0.3)
3	$V_{\text{ext}}$	PBE	0.1 (0.1)	0.1 (0.2)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
10	$V_{\text{ext}}$	PBE	0.1 (0.2)	0.2 (0.4)	0.4 (0.8)	0.2 (0.3)	0.8 (1.1)
1	$V_{\text{ext}}$	MGGA	0.1 (0.1)	0.1 (0.2)	0.1 (0.1)	0.1 (0.1)	0.2 (0.2)
2	$V_{\text{ext}}$	MGGA	0.1 (0.1)	0.2 (0.2)	0.1 (0.1)	0.0 (0.1)	0.2 (0.3)



TABLE II. (Continued.)

$N_{\text{electrons}}$	Input	Functional	$E_{\text{correlation}}$	$E_{\text{exchange}}$	$E_{\text{external}}$	$E_{\text{kinetic}}$	$E_{\text{total}}$
3	$V_{\text{ext}}$	MGGA	0.1 (0.1)	0.2 (0.3)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
10	$V_{\text{ext}}$	MGGA	0.1 (0.2)	0.3 (0.5)	0.3 (0.7)	0.2 (0.3)	0.7 (1.0)
1	$\rho$	LDA $\rightarrow$ PBE	0.1 (0.1)	0.1 (0.2)	0.1 (0.1)	0.1 (0.1)	0.2 (0.3)
2	$\rho$	LDA $\rightarrow$ PBE	0.1 (0.1)	0.2 (0.2)	0.1 (0.2)	0.1 (0.1)	0.3 (0.4)
3	$\rho$	LDA $\rightarrow$ PBE	0.1 (0.1)	0.1 (0.2)	0.2 (0.2)	0.1 (0.2)	0.4 (0.5)
10	$\rho$	LDA $\rightarrow$ PBE	0.1 (0.2)	0.2 (0.2)	0.3 (0.6)	0.2 (0.3)	0.8 (1.4)
1	$\rho$	LDA $\rightarrow$ MGGA	0.1 (0.1)	0.2 (0.3)	0.1 (0.2)	0.1 (0.2)	0.3 (0.3)
2	$\rho$	LDA $\rightarrow$ MGGA	0.1 (0.2)	0.3 (0.3)	0.1 (0.2)	0.1 (0.2)	0.3 (0.5)
3	$\rho$	LDA $\rightarrow$ MGGA	0.1 (0.2)	0.3 (0.4)	0.2 (0.3)	0.2 (0.2)	0.5 (0.7)
10	$\rho$	LDA $\rightarrow$ MGGA	0.1 (0.2)	0.2 (0.3)	0.4 (0.7)	0.2 (0.3)	0.9 (1.5)

## APPENDIX C: MEAN ABSOLUTE ERRORS WITH THE RND EXTERNAL POTENTIALS

TABLE III. Mean absolute errors (in mHa per electron) and root-mean-squared errors (in parentheses) for models trained in this report for the RND potentials. The abbreviations  $\rho$ ,  $V_{\text{ext}}$ , and predicted  $\rho$  are charge density, potential, and predicted charge density respectively. The arrows (i.e., LDA  $\rightarrow$  PBE) indicate that the charge density used as input to the DNN was calculated using the LDA exchange-correlation functional, but the labels (energies) were calculated using another exchange-correlation functional. The acronym DDE stands for density-driven error, as defined by Brockherde *et al.* [15]. The models labeled by  $V_{\text{ext}} \rightarrow \rho$  directly map the external potentials to charge densities. The models labeled by Predicted  $\rho$  map predicted charge densities to true energies.

$N_{\text{electrons}}$	Input	Functional	$E_{\text{correlation}}$	$E_{\text{exchange}}$	$E_{\text{external}}$	$E_{\text{kinetic}}$	$E_{\text{total}}$
1	$\rho$	LDA	0.9 (1.5)	1.4 (2.3)	14.5 (31.7)	2.5 (3.7)	14.5 (31.1)
2	$\rho$	LDA	1.1 (2.0)	1.8 (3.1)	9.6 (19.4)	2.2 (3.4)	9.9 (25.4)
3	$\rho$	LDA	2.0 (3.5)	3.1 (5.4)	35.1 (57.0)	5.6 (8.6)	36.1 (58.4)
10	$\rho$	LDA	2.0 (3.8)	3.0 (5.9)	73.4 (117.1)	6.7 (10.2)	74.4 (119.4)
1	$\rho$	PBE	1.9 (2.5)	3.1 (4.0)	15.3 (32.2)	3.8 (4.9)	15.1 (31.2)
2	$\rho$	PBE	1.3 (2.1)	1.9 (3.1)	9.5 (19.3)	2.2 (3.1)	10.1 (21.7)
3	$\rho$	PBE	2.0 (3.3)	2.9 (4.9)	34.4 (55.6)	5.2 (7.7)	35.6 (56.8)
10	$\rho$	PBE	1.9 (3.6)	2.8 (5.5)	73.7 (115.0)	7.6 (11.4)	75.0 (117.1)
1	$\rho$	MGGA	1.0 (1.7)	2.1 (3.4)	13.7 (30.7)	2.5 (3.8)	13.9 (31.0)
2	$\rho$	MGGA	1.3 (2.2)	2.6 (4.6)	10.0 (17.8)	2.4 (3.9)	10.2 (17.8)
3	$\rho$	MGGA	1.8 (2.8)	3.6 (5.7)	34.1 (55.7)	4.7 (7.1)	36.5 (62.2)
10	$\rho$	MGGA	2.2 (4.1)	4.4 (8.4)	77.2 (122.1)	7.3 (11.3)	78.5 (126.1)
1	$V_{\text{ext}}$	LDA	5.3 (9.0)	8.6 (14.4)	17.8 (28.7)	11.5 (19.2)	22.3 (36.8)
2	$V_{\text{ext}}$	LDA	5.3 (8.7)	8.5 (13.6)	18.5 (31.8)	8.8 (13.5)	21.2 (33.4)
3	$V_{\text{ext}}$	LDA	10.6 (15.0)	16.8 (23.7)	46.2 (70.2)	15.8 (23.6)	43.2 (66.7)
10	$V_{\text{ext}}$	LDA	6.6 (9.9)	10.3 (15.4)	50.3 (86.1)	10.5 (16.7)	40.9 (73.2)
1	$V_{\text{ext}}$	PBE	6.0 (9.9)	9.3 (15.2)	17.2 (27.3)	12.1 (19.3)	22.5 (35.6)
2	$V_{\text{ext}}$	PBE	6.0 (9.7)	9.0 (14.7)	19.5 (32.8)	10.0 (15.4)	21.3 (34.0)
3	$V_{\text{ext}}$	PBE	10.9 (15.3)	16.8 (23.8)	46.1 (69.5)	16.3 (24.2)	43.6 (65.4)
10	$V_{\text{ext}}$	PBE	7.1 (10.6)	10.8 (16.1)	49.3 (85.4)	10.4 (16.8)	39.4 (71.5)
1	$V_{\text{ext}}$	MGGA	5.7 (9.8)	12.0 (20.4)	16.1 (25.3)	10.9 (17.9)	21.7 (34.8)
2	$V_{\text{ext}}$	MGGA	5.4 (8.6)	11.0 (17.8)	14.8 (23.2)	8.1 (12.4)	19.3 (30.0)
3	$V_{\text{ext}}$	MGGA	11.1 (15.5)	22.9 (32.3)	44.5 (67.9)	16.7 (25.0)	43.0 (65.7)
10	$V_{\text{ext}}$	MGGA	7.3 (10.8)	15.0 (22.2)	50.8 (88.0)	10.8 (16.9)	41.4 (76.0)
1	$\rho$	LDA $\rightarrow$ PBE	1.2 (2.0)	1.7 (3.0)	14.5 (31.1)	2.4 (3.6)	14.6 (30.5)
2	$\rho$	LDA $\rightarrow$ PBE	1.2 (4.3)	1.8 (6.4)	8.6 (19.3)	1.9 (4.6)	9.1 (23.9)
3	$\rho$	LDA $\rightarrow$ PBE	1.9 (3.4)	2.8 (5.0)	34.8 (56.2)	5.1 (7.7)	35.6 (56.9)
10	$\rho$	LDA $\rightarrow$ PBE	2.2 (4.0)	3.2 (6.1)	75.5 (118.4)	7.6 (11.7)	76.7 (120.9)
1	$\rho$	LDA $\rightarrow$ MGGA	1.4 (2.8)	2.9 (5.8)	14.2 (30.9)	2.6 (4.5)	14.7 (31.6)
2	$\rho$	LDA $\rightarrow$ MGGA	1.5 (3.7)	3.0 (7.5)	8.0 (15.2)	2.0 (4.3)	8.7 (19.6)
3	$\rho$	LDA $\rightarrow$ MGGA	2.9 (5.4)	6.0 (11.3)	36.3 (58.2)	5.8 (10.0)	37.7 (59.8)

TABLE III. (Continued.)

$N_{\text{electrons}}$	Input	Functional	$E_{\text{correlation}}$	$E_{\text{exchange}}$	$E_{\text{external}}$	$E_{\text{kinetic}}$	$E_{\text{total}}$
10	$\rho$	LDA $\rightarrow$ MGGA	2.6 (4.9)	5.3 (9.9)	73.6 (115.6)	7.5 (11.7)	74.4 (117.4)
10	$V_{\text{ext}} \rightarrow \rho$	LDA	6.4 (11.0)	9.9 (17.1)	93.0 (151.4)	12.7 (21.5)	98.9 (167.5)
10	$V_{\text{ext}} \rightarrow \rho$ (DDE)	LDA	6.4 (10.7)	10.0 (16.7)	99.4 (154.4)	13.4 (22.2)	108.2 (175.5)
10	Predicted $\rho$	LDA	5.8 (10.1)	9.0 (15.7)	91.4 (143.2)	11.2 (19.1)	98.5 (158.7)
10	$V_{\text{ext}} \rightarrow \rho$	PBE	7.2 (11.7)	10.9 (17.7)	100.4 (164.8)	14.5 (24.6)	108.1 (185.8)
10	$V_{\text{ext}} \rightarrow \rho$ (DDE)	PBE	7.2 (11.5)	10.9 (17.7)	107.0 (168.1)	15.1 (24.9)	117.6 (193.7)
10	Predicted $\rho$	PBE	6.8 (11.1)	10.2 (16.8)	98.2 (151.5)	12.4 (21.4)	106.7 (170.3)
10	$V_{\text{ext}} \rightarrow \rho$	MGGA	7.5 (12.3)	15.4 (25.1)	94.6 (161.9)	13.0 (22.8)	103.3 (180.5)
10	$V_{\text{ext}} \rightarrow \rho$ (DDE)	MGGA	7.4 (12.0)	15.2 (24.6)	101.1 (162.4)	13.7 (23.3)	111.8 (184.0)
10	Predicted $\rho$	MGGA	7.1 (12.1)	14.7 (24.8)	94.9 (150.2)	12.7 (21.5)	102.5 (167.3)

## APPENDIX D: A NOTE ON THE DENSITY-DRIVEN ERRORS

In Table III, we report the mean absolute density-driven error rather than the density-driven error that is reported in [15]. When evaluating a machine-learning model, it is common to report absolute errors to avoid reporting an average error that would have error cancellation. Consider the total-energy expression

$$E[\rho] = F[\rho] + \int d\mathbf{r} V(\mathbf{r})\rho(\mathbf{r}), \quad (\text{D1})$$

where  $F$  is the universal functional defined in [1] and  $V$  is the external potential. The value  $E$  is the *true* energy given a *true* charge density  $\rho$ . The total error reported in [15] is defined to be

$$\Delta E = \tilde{E}[\tilde{\rho}] - E[\rho] = \Delta E_F + \Delta E_D, \quad (\text{D2})$$

where

$$\Delta E_F = \tilde{F}[\rho] - F[\rho] \quad (\text{D3})$$

is the functional-driven error and

$$\Delta E_D = \tilde{E}[\tilde{\rho}] - \tilde{E}[\rho] \quad (\text{D4})$$

is the density-driven error. The variables  $\tilde{\rho}$ ,  $\tilde{E}$ , and  $\tilde{F}$  represent a predicted charge density, an approximation to the total-energy functional, and an approximation to the true universal functional, respectively. If we consider the absolute value of the total error,

$$|\Delta E| = |\Delta E_F + \Delta E_D| \neq |\Delta E_F| + |\Delta E_D|, \quad (\text{D5})$$

then we can see that there must be error cancellation between the terms  $\Delta E_F + \Delta E_D$  in order for  $|\Delta E| < |\Delta E_D|$ . This is what we find in Table III.

- 
- [1] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).  
[2] J. Behler, *Phys. Chem. Chem. Phys.* **13**, 17930 (2011).  
[3] S. Lorenz, A. Groß, and M. Scheffler, *Chem. Phys. Lett.* **395**, 210 (2004).  
[4] C. M. Handley and P. L. A. Popelier, *J. Phys. Chem. A* **114**, 3371 (2010).  
[5] R. M. Balabin and E. I. Lomakina, *J. Chem. Phys.* **131**, 074104 (2009).  
[6] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).  
[7] C. Xie, X. Zhu, D. R. Yarkony, and H. Guo, *J. Chem. Phys.* **149**, 144107 (2018).  
[8] K. Schütt, P.-J. Kindermans, H. E. S. Felix, S. Chmiela, A. Tkatchenko, and K.-R. Müller, in *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017), pp. 991–1001.  
[9] K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, *Nat. Commun.* **8**, 13890 (2017).  
[10] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, *Phys. Rev. Lett.* **108**, 253002 (2012).  
[11] L. Li, T. E. Baker, S. R. White, and K. Burke, *Phys. Rev. B* **94**, 245129 (2016).  
[12] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, *Int. J. Quantum Chem.* **116**, 819 (2016).  
[13] J. C. Snyder, M. Rupp, K. Hansen, L. Blooston, K.-R. Müller, and K. Burke, *J. Chem. Phys.* **139**, 224104 (2013).  
[14] B. Kolb, L. C. Lentz, and A. M. Kolpak, *Sci. Rep.* **7**, 1192 (2017).  
[15] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, *Nat. Commun.* **8**, 872 (2017).  
[16] K. Ryczko, K. Mills, I. Luchak, C. Homenick, and I. Tamblyn, *Comput. Mater. Sci.* **149**, 134 (2018).  
[17] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).  
[18] K. Yao and J. Parkhill, *J. Chem. Theory Comput.* **12**, 1139 (2016).

- [19] M. S. Hybertsen and S. G. Louie, *Phys. Rev. B* **34**, 5390 (1986).
- [20] K. Mills, K. Ryczko, I. Luchak, A. Domurad, C. Beeler, and I. Tamblyn, *Chem. Sci.* **10**, 4129 (2019).
- [21] T. D. Kulkarni, W. F. Whitney, P. Kohli, and J. Tenenbaum, in *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., 2015), pp. 2539–2547.
- [22] K. Mills, M. Spanner, and I. Tamblyn, *Phys. Rev. A* **96**, 042113 (2017).
- [23] P. A. M. Dirac, *Math. Proc. Cambridge Philos. Soc.* **26**, 376 (1930).
- [24] M. F. Bloch, *Z. Phys. A Hadrons Nuclei* **57**, 545 (1929).
- [25] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [26] J. G. Vilhena, E. Räsänen, M. A. L. Marques, and S. Pittalis, *J. Chem. Theory Comput.* **10**, 1837 (2014).
- [27] S. Pittalis, E. Räsänen, N. Helbig, and E. K. U. Gross, *Phys. Rev. B* **76**, 235314 (2007).
- [28] C. Attaccalite, S. Moroni, P. Gori-Giorgi, and G. B. Bachelet, *Phys. Rev. Lett.* **88**, 256601 (2002).
- [29] M. A. Marques, M. J. Oliveira, and T. Burnus, *Comput. Phys. Commun.* **183**, 2272 (2012).
- [30] S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, *SoftwareX* **7**, 1 (2018).
- [31] X. Andrade, D. Strubbe, U. De Giovannini, A. H. Larsen, M. J. Oliveira, J. Alberdi-Rodriguez, A. Varas, I. Theophilou, N. Helbig, M. J. Verstraete *et al.*, *Phys. Chem. Chem. Phys.* **17**, 31371 (2015).
- [32] X. Andrade, J. Alberdi-Rodriguez, D. A. Strubbe, M. J. Oliveira, F. Nogueira, A. Castro, J. Muguerza, A. Arruabarrena, S. G. Louie, A. Aspuru-Guzik *et al.*, *J. Phys.: Condens. Matter* **24**, 233202 (2012).
- [33] See <http://clean.energyscience.ca/codes>.
- [34] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, Vol. 16 (USENIX Association, Berkeley, CA, 2016), pp. 265–283.
- [35] Y. Tang, [arXiv:1612.04251](https://arxiv.org/abs/1612.04251).
- [36] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [37] See <http://clean.energyscience.ca/datasets>.