
Machine learning and density functional theory

Ryan Pederson

Department of Physics and Astronomy, University of California, Irvine, CA 92697, USA

Bhupalee Kalita

Department of Chemistry, University of California, Irvine, CA 92697, USA

Kieron Burke*

*Department of Chemistry, University of California, Irvine, CA 92697, USA and
Department of Physics and Astronomy, University of California, Irvine, CA 92697, USA*

(Dated: May 22, 2022)

Over the past decade machine learning has made significant advances in approximating density functionals, but whether this signals the end of human-designed functionals remains to be seen. Ryan Pederson, Bhupalee Kalita and Kieron Burke discuss the rise of machine learning for functional design.

In 2016, DeepMind’s AlphaGo made history as the first computer program to defeat a professional human player in the board game Go. Aside from excelling at board games with human-designed rules, machine learning (ML) has seen impressive applications in scientific research [1] where the rules are instead governed by nature, such as protein folding and, more recently, density functional theory (DFT). But can the ML approach to DFT approximations significantly outperform human-designed ones and solve long-standing challenges?

1. DFT EVERYWHERE

In physical sciences, DFT is often the go-to computational method for solving electronic structure problems. DFT provides fully quantum solutions at a fraction of the cost of solving the Schrödinger equation directly, by mapping the coupled many-body problem to a single-particle problem. The electronic energy is considered as a functional (a function of a function) of the electron (probability) density, with only a small portion, the exchange-correlation energy, being approximated.

It is staggering to see just how important DFT calculations have become. Each year, tens of thousands of papers report useful predictions from DFT calculations, and today about one third of the National Energy Research Scientific Computing Center supercomputing resources use DFT [2]. John Perdew, who developed many of the formulas in current use, is one of the most cited physicists of all time.

2. THE GIGO PRINCIPLE IN DFT

The GIGO principle is an old adage in computing standing for: garbage in, garbage out. In DFT

this means that a calculation is only as good as the approximate functional used. Humans have worked at this for almost a century, and nowadays hundreds of different approximations are in use. Some build in well-studied limits, such as the uniform electron gas, and satisfy many known physical constraints of the exact functional, while others are tuned and fitted to reference datasets. Regardless, general failures have been identified over the years. A decade-old review [3] focused on the struggle to describe strongly correlated systems. This most grievous failure can be understood from the perspective of fractional charges (systems with noninteger total charge) and fractional spins (systems with noninteger spin magnetization). The exact energy is a linear interpolation of the energy of the adjacent integer systems, but approximations miss this, producing embarrassingly large systematic errors in strongly correlated systems as simple as stretched H_2 . Overcoming such fundamental DFT challenges is essential to expanding its applicability and reliability in condensed matter physics.

3. MACHINE LEARNING DFT

A proof of principle for ML-DFT appeared 10 years ago. For a simple problem, the kinetic energy of non-interacting fermions in a 1D box, a ML method (kernel ridge regression) could be used to find an approximation of the functional by training on examples from accurate numerical calculations [4]. The resulting functional was far more accurate than anything ever designed by humans, but only useful for simple model systems such as those it was trained on. The associated learning efficiency was also low, as hundreds of training examples were needed to reach high accuracy for a rather compact chemical space.

Later, the density was machine-learned directly from the external potential [5]. This demonstrated the practical usefulness of ML in DFT through realistic examples, such as proton transfer in a ML molecular

* kieron@uci.edu

dynamics simulation of malonaldehyde. However, unlike traditional DFT approximations, such ML models rarely generalize across elements.

Since then, there have been many attempts to bring the promise of ML to practical, generalizable functional construction. These efforts can be divided into two categories: those starting from traditional approximate forms suggested by humans (which are biased toward local and semi-local approximations) and those that use the entire density (that is, a non-local approximation) in some hard-to-understand way. Such non-local functionals can have poor generalizability, as the input feature space becomes vastly more complicated than local and semi-local forms which depend only on the density and its gradient at each point.

As described in Ref. [6] a neural network (NN) functional was trained on accurate densities as well as energies of just three molecules, producing semi-local ML approximations that worked as well as human-designed functionals for 150 test molecules, generalizing very well. A similar approach was used in Ref. [7] but non-local forms based on convolution NNs were also used to learn an entire dissociation curve within chemical accuracy, including the strongly correlated region, with only two training examples. The model also generalized well for other new (but similar) strongly correlated molecules that were not encountered in training. During training, an end-to-end differentiable DFT code (where all components are differentiable) was used to obtain gradient information by backpropagation through the entire self-consistent calculation. Such robust gradient-based training results in impressive generalization of functional approximations.

But the most recent exciting development comes once again from DeepMind [8]. A bevy of 17 researchers, using vast computational resources, revived an old human-designed suggestion, a local hybrid functional [9], that had been difficult to control. Their new NN-based functional, DM21, was trained by evaluating the energy non-self-consistently using approximate densities. The regression loss consisted of an energy loss plus an explicit gradient regularization term, thereby making this training approach substantially cheaper than Ref. [7]. DM21 was trained on thousands of molecular systems, orders of magnitude more than previous ML training sets, and outperforms most other

hybrid functionals on standard molecular benchmarks with impressive generalization. This ML functional can be used for main-group chemistry calculations, like most human-designed functionals. By including training on simple systems with fractional charges and spins, DM21 appears to perform significantly better than earlier approaches for strongly correlated systems. For instance, DM21 correctly dissociates systems such as H_2 , H_2^+ , and N_2 , meeting the long-standing DFT challenge of strong correlation in molecular systems.

4. WILL DFT GO THE WAY OF GO?

Researchers all over the world are currently trying out DM21, testing many different aspects to see if it lives up to its promise. The world of DFT applications is far too vast for DM21 developers to run even a fraction of useful tests in their original paper. Many promising approximations run into unexpected difficulties when tried in practice. The community will examine computational cost, accuracy, and transferability when testing DM21.

A basic issue is whether DeepMind’s approach can also work for materials. DM21 was trained and tested only on molecules, where wave function-based quantum chemistry provides high accuracy benchmark data. If such data were available for materials, for example, via computationally-expensive Quantum Monte Carlo simulations, a similar performance might be expected. However, the real aim of DFT is to find a single functional that works for both molecules and materials simultaneously, so that you can calculate everything in between, such as surfaces, clusters, and defects. Harsh experience suggests that accuracy for molecules degrades when good performance for solids is also required [10].

But, in the grand scheme of things, does this development signal the beginning of the end for human-designed functionals? Just as the world’s best Go player cannot compete with AlphaGo, can human insight and ingenuity long survive against huge data benchmarks, teams of coders, and almost unlimited computational resources? Or will human insight and ingenuity always be needed to choose the forms that machines learn from? The next decade will likely answer this question.

-
- [1] M. R. Douglas, *Nature Reviews Physics* **4**, 145 (2022).
 - [2] B. A. et al., Nersc-10 workload analysis (data from 2018), https://portal.nersc.gov/project/m888/nersc10/workload/N10_Workload_Analysis_latest.pdf (2020).
 - [3] A. J. Cohen, P. Mori-Sánchez, and W. Yang, *Chemical Reviews* **112**, 289 (2012).
 - [4] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, *Physical Review Letters* **108**, 253002 (2012).
 - [5] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman,

- K. Burke, and K.-R. Müller, *Nature Communications* **8**, 872 (2017).
- [6] R. Nagai, R. Akashi, and O. Sugino, *npj Computational Materials* **6**, 43 (2020).
- [7] L. Li, S. Hoyer, R. Pederson, R. Sun, E. D. Cubuk, P. Riley, K. Burke, *et al.*, *Physical Review Letters* **126**, 036401 (2021).
- [8] J. Kirkpatrick, B. McMorrow, D. H. Turban, A. L. Gaunt, J. S. Spencer, A. G. Matthews, A. Obika,

-
- L. Thiry, M. Fortunato, D. Pfau, *et al.*, *Science* **374**, 1385 (2021).
[9] F. G. Cruz, K.-C. Lam, and K. Burke, *The Journal of Physical Chemistry A* **102**, 4911 (1998).
[10] J. P. Perdew, *Science* **374**, 1322 (2021).

5. ACKNOWLEDGMENTS

Work supported by DOE DE-SC0008696 (R. P.) and NSF CHE-2154371 (B. K., K. B.).

6. COMPETING INTERESTS

The authors declare no competing interests.