

# Faraday Discussions

Accepted Manuscript



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

This article can be cited before page numbers have been issued, to do this please use: E. J. Landinez-Borda, K. O. Berard, A. Lopez and B. M. Rubenstein, *Faraday Discuss.*, 2024, DOI: 10.1039/D4FD00051J.

# Gaussian Processes for Finite Size Extrapolation of Many-Body Simulations

Edgar Josué Landínez Borda,<sup>1, a)</sup> Kenneth O. Berard,<sup>1</sup> Annette Lopez,<sup>2</sup> and Brenda Rubenstein<sup>1, b)</sup>

<sup>1)</sup>Department of Chemistry, Brown University, Providence, Rhode Island 02912, USA

<sup>2)</sup>Department of Physics, Brown University, Providence, Rhode Island 02912, USA

(Dated: 25 March 2024)

Key to being able to accurately model the properties of realistic materials is being able to predict their properties in the thermodynamic limit. Nevertheless, because most many-body electronic structure methods scale as a high-order polynomial, or even exponentially, with system size, directly simulating large systems in their thermodynamic limit rapidly becomes computationally intractable. As a result, researchers typically estimate the properties of large systems that approach the thermodynamic limit by extrapolating the properties of smaller, computationally-accessible systems based on relatively simple scaling expressions. In this work, we employ Gaussian processes to more accurately and efficiently extrapolate many-body simulations to their thermodynamic limit. We train our Gaussian processes on Smooth Overlap of Atomic Positions (SOAP) descriptors to extrapolate the energies of one-dimensional hydrogen chains obtained using two high-accuracy many-body methods: Coupled Cluster theory and Auxiliary Field Quantum Monte Carlo (AFQMC). In so doing, we show that Gaussian processes trained on relatively short, 10-30-atom chains can predict the energies of both homogeneous and inhomogeneous hydrogen chains in their thermodynamic limit with sub-milliHartree accuracy. Unlike standard scaling expressions, our GPR-based approach is highly generalizable given representative training data and is not dependent on systems' geometries or dimensionality. This work highlights the potential for machine learning to correct for the finite size effects that routinely complicate the interpretation of finite size many-body simulations.

Keywords: Machine Learning, Gaussian Process Regression (GPR), Gaussian Approximation Potentials, Smooth Overlap of Atomic Positions (SOAP) Descriptors, Coupled Cluster, Quantum Monte Carlo (QMC), Finite Size Extrapolation, Hydrogen Chains

## I. INTRODUCTION

Over the past few decades, *ab initio* electronic structure methods have transformed our ability to design materials by enabling researchers to predict the macroscopic and emergent behavior of solids from a basic knowledge of their constituent atoms. Researchers can now routinely model the electronic and geometric properties of systems ranging from quantum materials to heterogeneous catalysts with – or very near – chemical accuracy. However, the accuracy that accompanies many-body electronic structure methods such as Coupled Cluster (CC) theory, Quantum Monte Carlo (QMC), and many-body perturbation theories often comes at a steep cost: these methods typically scale as a high degree polynomial with system size. For example, Coupled Cluster Singles, Doubles, and Perturbative Triples [CCSD(T)] conventionally scales as  $O(N^3M^4)$ , where  $N$  is the number of electrons and  $M$  is the size of the basis set, while Auxiliary Field Quantum Monte Carlo (AFQMC) typically scales as  $O(N^2M^2 + M^2N)$ .<sup>1</sup> In contrast, mean field methods such as Density Functional Theory (DFT) scale as  $O(N^2 \log N)$ <sup>2,3</sup> or  $O(N)$ ,<sup>4</sup> when locality is a good approximation, but are only predictive when the degree of electron correlation is mild. Historically, the comparatively steep scaling of many-body methods has thwarted their direct application to solids with large unit and/or supercells, limiting their use to systems with just tens to, po-

tentially, hundreds of atoms. However, such smaller, more computationally-accessible systems cannot manifest the same long-range correlations as are present in larger, more realistic solids, and can exhibit spurious boundary effects that confound their interpretation. Indeed, given the remarkable accuracy of many modern electronic structure methods, these finite size errors are often the largest sources of error in many calculations of solids.<sup>5,6</sup> This leads to a long-standing conundrum: *if many-body methods may only be directly applied to smaller, finite systems, how can they be leveraged to predict the properties of larger, more realistic solids?*

To increase the feasibility of many-body methods for the prediction of the properties of solids in their infinite-size, “thermodynamic limit” researchers have developed approaches that correct results for smaller systems to predict the properties of larger systems. Such so-called finite size corrections consist of two main contributions: one-body and two-body corrections, which ameliorate the one- and two-body contributions to the total energy, respectively. One-body finite size errors typically stem from shell-filling effects that lead to a mis-estimation of the kinetic energy<sup>6,7</sup> and can therefore be corrected by a judicious averaging over k-points.<sup>8</sup> For example, in mean field theories, integrating over many points in the first Brillouin zone can be circumvented by instead approximating quantities using mean-value points known as Baldereschi points.<sup>9</sup> While many-body methods such as QMC methods need to integrate over the full simulation supercell, not just one point, twist averaging<sup>10</sup> provides a means of averaging over a set of angles (offset vectors) on the Brillouin zone of the supercell that results in a rapid convergence of the one-

<sup>a)</sup>Electronic mail: [edgar\\_landinez\\_borda@brown.edu](mailto:edgar_landinez_borda@brown.edu)

<sup>b)</sup>Electronic mail: [brenda\\_rubenstein@brown.edu](mailto:brenda_rubenstein@brown.edu)



body effects.<sup>6,10</sup> In contrast, two-body finite size effects stem from errors in the Coulomb and exchange-correlation interactions and are more challenging to correct. These effects can be alleviated by introducing modified versions of these interactions, such as model periodic Coulomb corrections.<sup>6,7</sup> An alternative approach for correcting both one- and two-body finite size effects is to determine finite size corrections using methods that scale more gracefully with system size, such as Density Functional Theory (DFT).<sup>6,7,11</sup> Such methods are used to estimate the differences in energies between smaller and larger systems, and then these differences are added to the smaller-sized many-body calculations. One such DFT-based approach is the Kwee, Zhang, Krakauer correction (KZK).<sup>12</sup> While such corrections are now widely applied to materials, they inherently lack the accuracy that would be possible if many-body corrections that take strong correlation into account were applied.

Even though these one- and two-body corrections markedly reduce finite-size errors, extrapolations to the thermodynamic limit are often still made to reduce any remaining errors. The simplest approach for performing these extrapolations is to fit many-body results obtained at smaller system sizes (e.g., 2x2 or 3x3 supercells) to functional forms that enable extrapolation to larger system sizes.<sup>7,11</sup> Nevertheless, it is often unclear which functional form should be employed since it can vary with the geometry, dimensionality, and electronic phase of the material.<sup>5,8</sup> This is especially true for systems with atypical geometries and boundary conditions. It also of particular importance for calculations involving excited states, including gap and exciton binding energy calculations, because excited states can be more difficult to converge to their thermodynamic limit.<sup>13–15</sup> When a system's correlation energy converges slowly, the number of points necessary for accurate fitting can exceed computational constraints, limiting the overall utility of such extrapolations and the results they yield.<sup>8,16</sup>

One potentially promising approach for estimating many-body corrections that can reduce this computational expense is machine learning. Machine learning methods surrogate more complex models with regressions that have lower computational complexity, thereby accelerating prediction.<sup>17–19</sup> In the context of condensed matter physics, machine learning has been employed to accelerate the prediction and discovery of new materials based upon the properties of known materials<sup>20</sup> as well as to learn the presence of certain phases based upon their known signatures.<sup>21,22</sup> Machine learning techniques have moreover recently been harnessed to accelerate and improve the accuracy of quantum Monte Carlo methods (see a more detailed discussion in Section II).<sup>8,18,23–28</sup> To approach the problem of determining accurate, many-body finite size corrections, one can analogously imagine using data from smaller system sizes to train machine learning algorithms to predict the properties of systems of larger sizes. An early such work used energies and densities from the density matrix renormalization group to learn the DFT kinetic energy functional of hydrogen chains in the thermodynamic limit.<sup>29</sup> More recently, while this work was being prepared, Gaussian process regression techniques were shown to be able to successfully learn corrections to coupled cluster calculations in k-space. More

specifically, Mihm *et al.*<sup>8,30</sup> employed the transfer structure factor to quantify the finite size effects present in coupled cluster theories' correlation energy. They then innovatively bypassed directly computing the structure factor for G values approaching zero (i.e., in the thermodynamic limit) by flexibly representing the structure factor using Gaussian Process Regression.

In this work, we leverage Gaussian Process Regression (GPR)<sup>31</sup> to learn finite size corrections in real-space to homogeneous (one-dimensional) and inhomogeneous (two-dimensional) hydrogen chains modeled using the first-principles, many-body methods Coupled Cluster (CC) Theory and Auxiliary Field Quantum Monte Carlo (AFQMC). Kernel methods like Gaussian processes<sup>31</sup> are advantageous because they are not parametric and make use of Bayesian inference that can come at a lower  $O(N_t^3)$  (where  $N_t$  is the size of the training set) cost than more complicated parametric methods such as neural networks that scale with the number of layers employed.<sup>17</sup> Gaussian processes have also been shown to make equally, if not more, accurate predictions than neural networks when less training data is available, which is an important consideration when training is to be performed on data generated using relatively expensive electronic structure calculations.<sup>32</sup> We use Gaussian processes to first predict the energies of one-dimensional, homogeneous hydrogen chains of varying lengths using atomic environment descriptors that enable us to incorporate information regarding the geometry and electronic density of each atom and its neighbors. Importantly, even though machine learning methods are most accurate for interpolation, we demonstrate that training our models on the energies of one-dimensional hydrogen chains containing 10-30 atoms enables us to predict (extrapolate) the energies of chains of more than 100 atoms, nearing the thermodynamic limit, with sub-milliHartree accuracy. To contextualize the accuracy of our methods, we compare the accuracy of our predictions to that of polynomial fits to larger-sized systems, the so-called “subtraction trick,”<sup>33</sup> and other alternative regression methods. Finally, to demonstrate the generalizability and robustness of our approach, we show that our technique can readily be adapted to also extrapolate the energies of heterogeneous chains of hydrogen dimers, which possess more free parameters, to their thermodynamic limit. This work thus illustrates that machine learning is a relatively cheap, yet accurate means of correcting for finite size effects in many-body simulations that can potentially address many of the challenges the many-body modeling community faces predicting the properties of solids in the thermodynamic limit.

In the spirit of a *Faraday Discussion*, in Section II, we begin with a discussion of the emerging synergies between machine learning techniques and stochastic electronic structure methods. We then describe the machine learning methods, descriptors, and electronic structure techniques we employ in our finite-size extrapolation research in Section III. We next present our primary results demonstrating our technique's ability to accurately correct for finite size errors in Section IV. We conclude by discussing the relative merits and potential applications of our algorithm in Sections V and VI.



## II. MACHINE LEARNING IN STOCHASTIC ELECTRONIC STRUCTURE

Over the past decade, an increasing amount of research has shown that stochastic electronic structure and machine learning methods can form a very fruitful partnership that both accelerates and extends the capabilities of stochastic methods. Because stochastic electronic structure methods are often more expensive than other common electronic structure methods such as Density Functional Theory, machine learning techniques hold the promise of making stochastic electronic structure techniques less costly. At the same time, the high accuracy of most stochastic electronic structure techniques like Diffusion,<sup>34,35</sup> Full Configuration Interaction,<sup>36</sup> and Auxiliary Field<sup>37,38</sup> Quantum Monte Carlo methods can provide ML techniques with high-quality data that can be used to correct less accurate predictions.

One triumph of the union of these techniques has been the generation of machine learned force fields from QMC energies and gradients.<sup>27,39–41</sup> QMC energies and forces calculated for representative configurations are used to train a variety of different neural networks, e.g., Behler-Parrinello Neural Networks,<sup>42</sup> or other architectures, which are in turn used to predict the energies and forces for other configurations, accelerating geometry relaxation and/or *ab initio* molecular dynamics simulations.<sup>43,44</sup> For instance, in recent work, Diffusion Monte Carlo energies and forces were used to generate a force field using a hierarchical  $\Delta$ -machine learning scheme based upon the Deep Potential Molecular Dynamics (DPMD) framework<sup>45</sup> that was able to successfully uncover a new phase of hydrogen.<sup>27</sup> Since QMC has historically met challenges calculating forces,<sup>46</sup> recent work has also exploited machine learning architectures to learn force fields from *energy data alone*.<sup>39</sup> Other ways to further reduce the cost of QMC data generation for training itself employ either  $\Delta$ -ML<sup>47</sup> or transfer learning.<sup>48</sup> These techniques first learn potentials and forces, using data from less accurate, but less costly theories and then correct those force fields by either adding a machine learned correction or updating the less accurate force field with select higher accuracy information. Both methods capitalize on the fact that less accurate theories can often reproduce much of the correct physical behavior of a system, meaning that high accuracy methods are effectively only needed to correct specific phenomena or regions of the potential energy surface. Further opportunities lie in better harnessing the statistical nature of stochastic methods to more efficiently train such force fields.<sup>49</sup> Overall, QMC-quality force fields open up the grand possibilities of studying dynamics in large molecular or solid state systems with relatively little overhead, making QMC dynamics a practical reality.

Stochastic methods and machine learning techniques have also been fruitfully paired to develop new neural network-based variational ansätze. The Variational Principle, which states that the ground state wave function of a system can best be approximated by varying the parameters and forms of trial wave functions to minimize the energy of the system, has long been used to produce wave function ansätze in computational quantum chemistry and physics. Often, such ansätze have

been optimized using QMC (i.e., Variational Monte Carlo methods) and used either on their own or as starting points for projection-based QMC techniques.<sup>6,50</sup> Historically, the forms of these variational ansätze have been specified based upon knowledge of the chemistry/physics they ultimately aim to describe (e.g., Gutzwiller<sup>51</sup> or pairing<sup>52</sup> wave functions) or confidence that their form is generalizable and expressive enough to describe the phenomenon under study (e.g., backflow wave functions).<sup>53</sup> Specifying the forms of trial wave functions based upon the physics expected can potentially lead to circular logic in which the physics that is expected to be seen is incorporated into a variational wave function form that then recovers that physics.

Recently, machine learning has been employed to overcome this limitation by providing a means of creating highly expressive variational wave functions. One of the most popular means of achieving this has been to use deep neural networks to specify a given variational wave function and then to optimize that neural network using the energy and/or variance as a loss function.<sup>54</sup> Examples of such variational neural networks include DeepQMC,<sup>55</sup> FermiNet<sup>23</sup>, and PauliNet,<sup>24</sup> all of which have shown promise determining the ground states of challenging chemical systems. PauliNet and FermiNet, for example, use deep neural networks to learn a parameterized form of the Jastrow factor and backflow functions and maintain antisymmetry using Slater determinants. Unlike traditional methods that use single-particle orbitals, FermiNet employs functions invariant under two-electron permutations and incorporates back-flow-like transformations for enhanced accuracy.<sup>56</sup>

An alternative approach to combining the expression and optimization of wave functions with machine learning has been neural network quantum states.<sup>57</sup> One promising form of neural network quantum states established by Carleo and Troyer are Restricted Boltzmann Machines, which implement a representation of the wave function through hidden and visible layers.<sup>57</sup> The Boltzmann distribution models the probabilities associated with different configurations of visible and hidden nodes based on the energy; lower energies are favored to accommodate the variational principle which guides the optimization of wave function parameters. These wave functions can then be extrapolated to larger systems by reusing the learned features of the wave function to initialize a machine learning model applied to a similar, but larger system.<sup>58,59</sup> This process of transferring the learning done for one type of problem to a related, but different problem makes seemingly out-of-reach problems, such as the thermodynamic limit, computationally feasible. Success with the transverse-field Ising model,<sup>57</sup> Heisenberg model,<sup>57</sup> and molecules<sup>60</sup> has been demonstrated. Akin to the use of GPR in this work, Gaussian Processes have also been used to specify wave functions called Gaussian Process States.<sup>61</sup> These wave functions are expressed as the exponential of a GP estimator and thus, as Gaussian processes more generally, are highly generalizable and can provide critical information about uncertainties. Such machine learning-based wave functions offer a potential means of achieving unprecedented levels of accuracy without the need for typically more expensive projection techniques.





Given these successes combining stochastic methods with machine learning approaches - and the many more we have not been able to discuss due to space constraints - here, we focus on the possibility of using machine learning methods to extend QMC's capabilities in a different way: by facilitating the extrapolation of QMC results to the thermodynamic limit.

### III. METHODS

#### A. Gaussian Approximation Potentials (GAP)

In this work, we employ Gaussian Process Regression (GPR) to predict finite size corrections for discrete hydrogen chains. We have focused on GPR<sup>62</sup> because it has previously been shown to yield high accuracy results with less training data than comparable methods.<sup>32</sup> This is an especially desirable property when one is interested in performing regressions on data obtained from comparatively costly many-body simulations, since computational expense practically limits how much reference data can reasonably be collected. The Bayesian nature of GPR also makes it possible to compute the variance of its predictions, which greatly facilitates the interpretation of its results.<sup>63</sup> For these reasons, we employ a GPR-based approach which is very similar in flavor to the Gaussian Approximation Potential (GAP) approach.<sup>64</sup> We first summarize our approach at a high level and then provide more details in subsequent subsections.

A GPR is a random process which takes input vectors  $\mathbf{x}_i$  and maps them to random variables  $y = f(\mathbf{x})$  with a multivariate, normal joint distribution with covariance ( $K$ )<sup>31,63</sup>

$$p(f(x)) \sim N(\mu, K). \quad (1)$$

The target function  $f(\mathbf{x}_i)$  (which yields the energy in this work) is characterized by the expectation value of the distribution  $\mu = \langle f(\mathbf{x}_i) \rangle$ . Like GAP, we use atomic environment descriptors<sup>65</sup> as input features,  $\mathbf{x}_i$  (which are vectors containing the atomic descriptors of a structure  $i$ ). These capture the main features of the electron density of an atom and its neighborhood (its atomic environment) to represent the electronic characteristics of the atoms. The covariance determines how the features are correlated and is specified by the kernel function. In kernel methods such as GPR,<sup>31,66</sup> input features  $\mathbf{x}_i$  are mapped to a nonlinear, high-dimensional space through the function  $\phi(\mathbf{x}_i)$ . Correlations between descriptors that represent different atomic structures are subsequently represented by taking their inner product in this nonlinear space to yield the kernel

$$K(x_i, x_j) = \phi(\mathbf{x}_i) \cdot \phi(\mathbf{x}_j). \quad (2)$$

Nevertheless, the kernel can be defined in a more arbitrary way as long as it satisfies the properties of a covariance matrix.<sup>66</sup> In order to make predictions, Bayesian inference can be used to compute new values of the target function.<sup>31,63</sup> This is done by extending the distribution to unobserved data,  $\mathbf{y}^*$ .

The idea is to generate a distribution based on the observed data ( $\mathbf{y}, \mathbf{X}$ ) using unseen data  $\mathbf{X}^*$  to generate the prediction  $\mathbf{y}^*$  with the corresponding joint distribution:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}^* \end{bmatrix} \sim N \left( \begin{bmatrix} \mu \\ \mu^* \end{bmatrix}, \begin{bmatrix} K & K_* \\ K_*^T & K_{**} \end{bmatrix} \right), \quad (3)$$

where  $\mu$  and  $\mu_*$  denote the means over the training and unobserved data, respectively, and  $K$ ,  $K_*$ , and  $K_{**}$  represent the covariances among the training data, training and unobserved data, and unobserved data, respectively. Based upon Bayes' Rule, the posterior distribution is Gaussian since the joint distribution is Gaussian. The posterior distribution can be expressed as

$$P(\mathbf{y}^* | \mathbf{y}) \sim N(\hat{y}, \hat{K}), \quad (4)$$

while the predicted mean and variance for an unobserved point may be expressed as

$$\mathbf{y}^* = \hat{y} = K^{-1} \mathbf{y} K(\mathbf{X}, \mathbf{X}^*) \quad (5)$$

and

$$\sigma^* = K(\mathbf{X}^*, \mathbf{X}^*) - K(\mathbf{X}, \mathbf{X}^*)^T K(\mathbf{X}, \mathbf{X})^{-1} K(\mathbf{X}, \mathbf{X}^*). \quad (6)$$

The functional form of the prediction is equivalent to that produced by Kernel Ridge regression,<sup>66</sup> and can be written in the same way

$$\mathbf{y}^* = \hat{y} = (K + \sigma^2)^{-1} \mathbf{y} K(\mathbf{X}, \mathbf{X}^*). \quad (7)$$

In this equation, the weights,  $\alpha$ ,<sup>66</sup> are given by

$$\alpha = (K + \sigma^2 I)^{-1} \mathbf{y}. \quad (8)$$

Equation 7 can be written in terms of the coefficients given by Equation 8

$$\mathbf{y}^* = \sum_i \alpha_i \cdot K(\mathbf{x}_i, \mathbf{x}^*), \quad (9)$$

where the  $\alpha_i$  are vectors of the coefficients obtained from the regression and  $K(\mathbf{x}_i, \mathbf{x}^*)$  is the kernel between the unseen data,  $\mathbf{x}^*$ , and the training data,  $\mathbf{x}_i$ . Kernel methods such as GPR can thus be used to predict the total energy,  $E_{total}^*$ , using the equation

$$E_{total}^* = \sum_i \alpha_i \cdot K(\mathbf{x}_i, \mathbf{x}^*). \quad (10)$$

The kernels can be tuned to optimize the prediction of the Gaussian process through the selection of their free parameters, known as hyper-parameters. The most common method of optimizing the posterior is the log-likelihood maximization method. In this work, we use three-way hold-out and log-likelihood maximization over the hyper-parameters.



## B. Atomic environment descriptors and regression model

In contrast with physics-based approaches for describing a system, machine learning models are often more expressive, meaning that a single model has the potential to describe many different systems. One way to constrain the predictions of a machine learning model is to include prior physical information in the surrogate model. This can be achieved by making the model invariant to symmetries, including translational, rotational, or permutation symmetries, or constraints, in order to suppress spurious correlations. These symmetries or constraints are usually incorporated into the model in two ways: explicitly integrating these symmetries into the regression algorithm or designing features that are invariant to the symmetry transformations.

Here, we incorporate symmetries via the latter approach using Smooth Overlap of Atomic Positions (SOAP) descriptors that are invariant to rotation and translation. These atomic environment descriptors represent the electron density at some point  $r$  by the superposition of the Gaussian densities of atoms with the same atomic number  $Z$  in the neighborhood of that point

$$\rho^Z(r) = \sum_i^{\|Z_i\|} \exp\left(-\frac{\|r - R_i\|^2}{2\sigma^2}\right), \quad (11)$$

where  $R_i$  is the position of an atom,  $i$ , in the neighborhood and  $\sigma^2$  is the variance of the Gaussian. This density may be expanded in terms of radial and angular basis functions

$$\rho^Z(r) = \sum_{nlm} c_{nlm}^Z Y_{lm} g_n(r), \quad (12)$$

where the  $g_n(r)$  are the  $n$  radial basis functions that can be expressed in terms of polynomials or atomic orbitals and the  $Y_{lm}$  correspond to the spherical harmonic functions. The  $c_{nlm}^Z$  coefficients of the expansion can be computed by integrating over the density

$$c_{nlm}^Z(\mathbf{r}) = \iiint_{\mathcal{R}^3} dV g_n(r) Y_{lm}(\theta, \phi) \rho^Z(\mathbf{r}). \quad (13)$$

In this work, we use the Dscribe library<sup>67</sup> to obtain the descriptors. This library implements SOAP descriptors using a partial power spectrum that only includes real spherical harmonics. Because the density depends on the square of the distances between points, it is already invariant to translation. A descriptor vector,  $\mathbf{p}$ , is formed from elements of the power spectrum

$$p(\mathbf{r})_{nr'l}^{Z_1 Z_2} = \pi \sqrt{\frac{8}{2l+1}} \sum_m c_{nlm}^{Z_1}(\mathbf{r})^* c_{n'l m}^{Z_2}(\mathbf{r}), \quad (14)$$

where  $n$  and  $n' \leq n_{max}$  run over the radial basis functions and  $l \leq l_{max}$  runs over the spherical harmonics.  $n_{max}$  and  $l_{max}$  define the maximum number of radial and angular functions in which the density in Equation 12 is expanded, respectively.  $Z_1$

and  $Z_2$  are the atomic numbers of the species. The resulting power spectra are rotationally- and permutationally-invariant by construction.

The original SOAP descriptors compare the local atomic environments using a kernel that is the dot product of the normalized power spectra between different configurations

$$K^{SOAP}(\mathbf{p}, \mathbf{p}') = \left( \frac{\mathbf{p} \cdot \mathbf{p}'}{\sqrt{(\mathbf{p} \cdot \mathbf{p}) (\mathbf{p}' \cdot \mathbf{p}')}} \right)^\xi. \quad (15)$$

This kernel takes the overlap of two atomic environments. However, other kernels employ different ways of measuring the similarity of the environments that may lead to better results. One of the most common kernels because of its versatility and robustness is the Radial Basis Function (RBF) or Squared Exponential (SE) kernel

$$K(\mathbf{p}, \mathbf{p}') = v^2 \exp\left(-\frac{d(\mathbf{p}, \mathbf{p}')^2}{l^2}\right), \quad (16)$$

where  $d(\mathbf{p}, \mathbf{p}')$  is the Euclidean distance,  $v^2$  is a tunable amplitude, and  $l^2$  is the global weight or length scale of the features. We choose to use the latter kernel throughout this work because of its flexibility and robustness for comparing features.

## C. Comparing environments with global descriptors

The descriptor vector,  $\mathbf{p}$ , of an atomic structure depends on the number of atoms of each species and is created by concatenating the different combinations of atomic species, each with  $n$  radial basis functions and a maximum angular number  $l_{max}$ .<sup>65</sup> As a result, structures with different numbers of atoms,  $M$ ,  $N$ , have different numbers of descriptors. One way to deal with descriptor vectors of differing lengths is to pad the feature vectors with zeros such that their dimensions match those of the descriptor vectors with the largest number of features in the samples. A similar approach involves padding the dummy (missing) features with values selected to decrease the biases the missing features would otherwise introduce.<sup>68</sup>

An alternative that can reduce bias is the use of global descriptors. These descriptors characterize the whole structure, i.e., the features depend on all of the atoms, rendering the number of features independent of the number of atoms in the structure. However, such an approach may diminish the quality of the kernel, since the descriptors may not have enough resolution to distinguish subtle differences between structures because of their global nature. A very simple and intuitive method to make the kernel global is to construct an "average kernel."<sup>69</sup>

$$K(A, B) = \frac{1}{NM} \sum_{i,j}^{N,M} C_{ij}(A, B). \quad (17)$$



Such a kernel recursively compares the features of the atoms  $i$  and  $j$  in structures  $A$  and  $B$ , respectively, using the kernel,  $C$ , and averaging over its corresponding numbers of atoms  $N$  and  $M$ . This approach is equivalent to averaging the features of all of the atoms of each configuration and comparing them with the kernel  $C$ , which amounts to making the descriptors global

$$\bar{p}(\mathbf{r})_{nm'l}^{Z_1 Z_2} = \pi \sqrt{\frac{8}{2l+1}} \sum_i \frac{1}{N} \sum_m \left( c_{nlm}^{i,Z_1}(\mathbf{r})^* \right) \left( c_{n'lm}^{i,Z_2}(\mathbf{r}) \right). \quad (18)$$

The RBF kernel with the global descriptors then becomes

$$K(\bar{p}, \bar{p}') = v^2 \exp \left( \sum_i \frac{d(\bar{p}_i, \bar{p}'_i)^2}{l_i^2} \right). \quad (19)$$

It is important to note that, when global descriptors are employed, the total energy is no longer the simple sum of local contributions. Now, it explicitly depends on quantities that interrelate features of the whole structure. This overall description of atomic structures implicitly removes the need for descriptors that capture long-range order. Nonetheless, the resolution of the features still needs to be high enough to capture small structural changes, as mentioned earlier. The resolution of the kernel can be improved by weighting each global feature by some characteristic length,  $l_i$ , according to Equation 19. This improves kernel performance by allowing fine-tuning of the parameters, but at the cost of adding more complexity to the model. In the following, we employ this combination of SOAP-averaged descriptors and the RBF kernel on linear hydrogen chains, which serve as an interesting and challenging benchmark.

## IV. RESULTS

### A. One-Dimensional, Homogeneous Hydrogen Chains

#### 1. Coupled Cluster and AFQMC Database of Homogenous Hydrogen Chain Energies

To analyze the ability of our GPRs to predict the energies of solids in their thermodynamic limit, we first attempt to predict the finite size effects of linear hydrogen chains (LHC) stretched homogeneously, i.e., with their atoms equally-spaced, and with open boundary conditions. This system is a very well-known benchmark for strong electron correlation because of the multireference character it develops at long bond distances and has therefore been used to test the accuracy of a wide-range of many-body methods.<sup>1,70-73</sup> As illustrated in Figure 1, Unrestricted Hartree-Fock theory (UHF) underbinds the hydrogen atoms, while Unrestricted Coupled Cluster Theory (UCCSD(T)) and AFQMC are able to relatively accurately reproduce the chains' energies near their equilibrium bond lengths, but can struggle to capture their energies at longer bond lengths closer to the dissociation limit.<sup>1,70</sup>

This system furthermore exhibits a metal-to-insulator transition when stretched homogeneously, which occurs at 1.8 bohr.<sup>1</sup> This transition is of second order, meaning that it is continuous with respect to the energy, but can be characterized by the polarization or spin correlation functions.<sup>70</sup> Dimerization of pairs of hydrogen atoms in the chains can be observed by looking at the electron density profile along the chains, as in Figure 2. The maxima correspond to the nuclear positions, while the deep minima indicative of dimerization may be observed between pairs of atoms at all of the chain lengths depicted. Methods capable of predicting the energies as a function of bond length must implicitly be able to predict energies across these transitions.

In order to generate enough data for training, we created a database of the energies of hydrogen chains at varying bond lengths using UHF and two many-body methods - UCCSD(T)<sup>74</sup> and AFQMC<sup>75</sup> - in the minimal STO-6G basis. UCCSD(T) has long been considered the gold standard for accuracy for quantum chemistry calculations,<sup>74,76</sup> and is seeing an increasing number of applications to solids.<sup>77,78</sup> AFQMC<sup>75</sup> is a second-quantized QMC method that, despite its typical use of the phaseless approximation,<sup>79</sup> has been shown to achieve chemical accuracy in systems ranging from small molecules,<sup>80-82</sup> to complexes,<sup>82,83</sup> to strongly correlated solids.<sup>84,85</sup> As a check on our databases, we produced and extended the benchmarks of Motta *et al.*<sup>1</sup> with sub-milliHartree accuracy (see Figure 1).

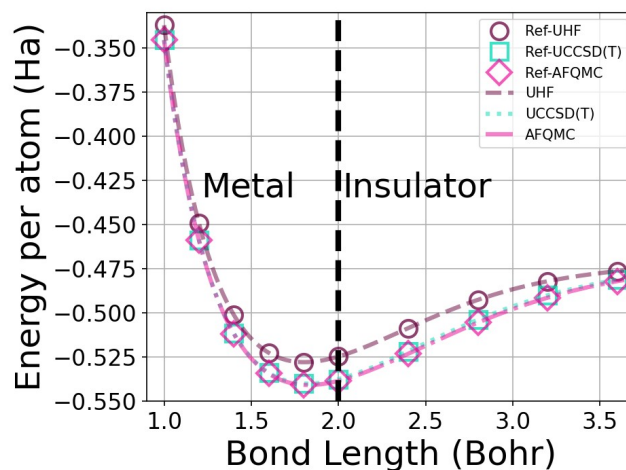


FIG. 1. Energy per atom vs. bond length for a 50-atom hydrogen chain using the UHF, UCCSD(T), and AFQMC methods in the STO-6G basis. The symbols depict the energies from calculations from Reference ('Ref') 1, while the dotted lines interpolate among 250 of our database energies. AFQMC error bars are too small to see.

To perform our UHF and UCCSD(T) calculations, we use the open source software PySCF.<sup>86</sup> For the AFQMC calculations, we use the high-performance implementation of AFQMC in QMCPACK.<sup>87</sup> Within QMCPACK, we employ UHF wave functions produced by PySCF as trial wave functions and perform calculations with a time step of 0.005, 1000 walkers, a Cholesky decomposition threshold of  $10^{-8}$ , and  $10^4$  steps in the phaseless approximation.<sup>75</sup> Energies






FIG. 2. Electron density as a function of atomic position ( $x$ ) for 10-atom hydrogen chains. Each curve depicts the electron density profile when the chain is stretched homogeneously at the bond lengths indicated in the legend. The change in the distance between and depth of adjacent local minima as the bond distance is increased reflects the onset of dimerization. The densities depicted here were computed using Full-Configuration Interaction<sup>74</sup> and the y-axis was shifted so that the profiles for all bond lengths could be clearly seen.

are computed with the hybrid estimator. Using all of these methods, we compute 250 points for each 10-60-atom chain with bond lengths ranging from 1 to 3.65 bohr. For chains of 70 to 100 atoms, we compute 40 points within the same range of bond lengths in order to conserve computational resources. The 10-30 atom data was used for training, while chains with larger numbers of atoms were used for benchmarking and analysis.

## 2. Energy Predictions Using Gaussian Process Regression

We use the smallest of our hydrogen chains of 10-30 atoms to train and test the GP regressions, which corresponds to 750 total samples. Samples were uniformly mixed by shuffling the data points at all bond lengths for each chain of a given size in the training set. This is to avoid training with an imbalanced data set. SOAP descriptors were constructed by using six GTOs as radial basis functions with a sigma of 1 bohr and six tesseral spherical harmonics as angular functions to build the atomic descriptors for all sizes and bond lengths. A cutoff radius that defines the extent of the atomic environment was set to 7 bohr for all chain sizes and bond lengths. This cutoff radius guarantees that the local environment of an atom consists of a maximum of 14 atoms at the shortest bond lengths studied and a minimum of 2 atoms at the longest bond lengths studied. It may be anticipated that the local atomic environment descriptors become linearly dependent when they have a large cutoff radius and are placed on bulk atoms that repeat throughout the chains. Nonetheless, descriptors placed on the edge atoms manifest asymmetries that reflect the finite extent

of the chains.

The descriptors are first generated for all of the atoms of each chain in the database. A global descriptor is then obtained by averaging each descriptor over the atoms within each chain. Finally, feature selection is carried out by obtaining leverage scores from a CUR decomposition.<sup>88</sup> The leverage scores are ordered in descending order and features are taken until 97% of the leverage score is accounted for. To perform the CUR decomposition, a Singular Value Decomposition (SVD)<sup>88</sup> is conducted given a singular value threshold that defines the rank of the decomposition. For this purpose, we used optimal hard thresholding,<sup>89</sup> which makes an optimal choice based on the dimensions and the estimated noise in the features or global descriptors matrix. We don't orthogonalize the features or use covariate principal coordinate analysis to improve our current feature selection, as further discussed in Section V.<sup>90</sup>

A Gaussian kernel with multiple length scales allows more sensitivity to global descriptors without greatly increasing the complexity of the model. We used the maximum likelihood<sup>31</sup> method to optimize the kernel hyper-parameters. We employed up to 500 configurations for training and 250 for validation.

## 3. Accuracy of GPR Predictions

After training our GPRs on the UCCSD(T) and AFQMC energies of shorter hydrogen chains, we are able to predict the energies per atom of chains with larger numbers of hydrogen atoms over the same range of bond lengths in the database with reasonable accuracy. We predict the energies per atom using the mean and variance of the posterior distribution.

Figure 3 depicts the differences between the energies computed with the UCCSD(T) (left) and AFQMC (right) methods, and their respective GPR predictions. In both cases, the differences between the predictions and the calculated energies are less than 1 mHa. It is reassuring to note that the short chain length predictions are most accurate throughout the prediction interval, which is a consequence of training the Gaussian processes on short chains. Prediction errors grow with the lengths of the chains because the generalization error increases with system size. This is reflected in the larger confidence intervals that accompany the larger chain length predictions. Hydrogen chains have previously been observed to exhibit slower convergence at short bond lengths because their total chain lengths are not yet long enough to converge finite size effects that stem from long-range Coulomb interactions. This comparatively slow convergence is likely responsible for the larger error bars we observe at short bond lengths.<sup>1</sup> At bond lengths longer than 3 Bohr where dissociation begins to occur, the error is significantly smaller and expected to converge faster because chains with longer total lengths will more rapidly converge the long-range Coulomb interaction.

The quality and characteristics of the AFQMC-based GPR predictions are similar to that of the UCCSD(T)-based predictions. Higher accuracies are again observed for shorter chains and at larger bond lengths. The AFQMC-GPR differences are,





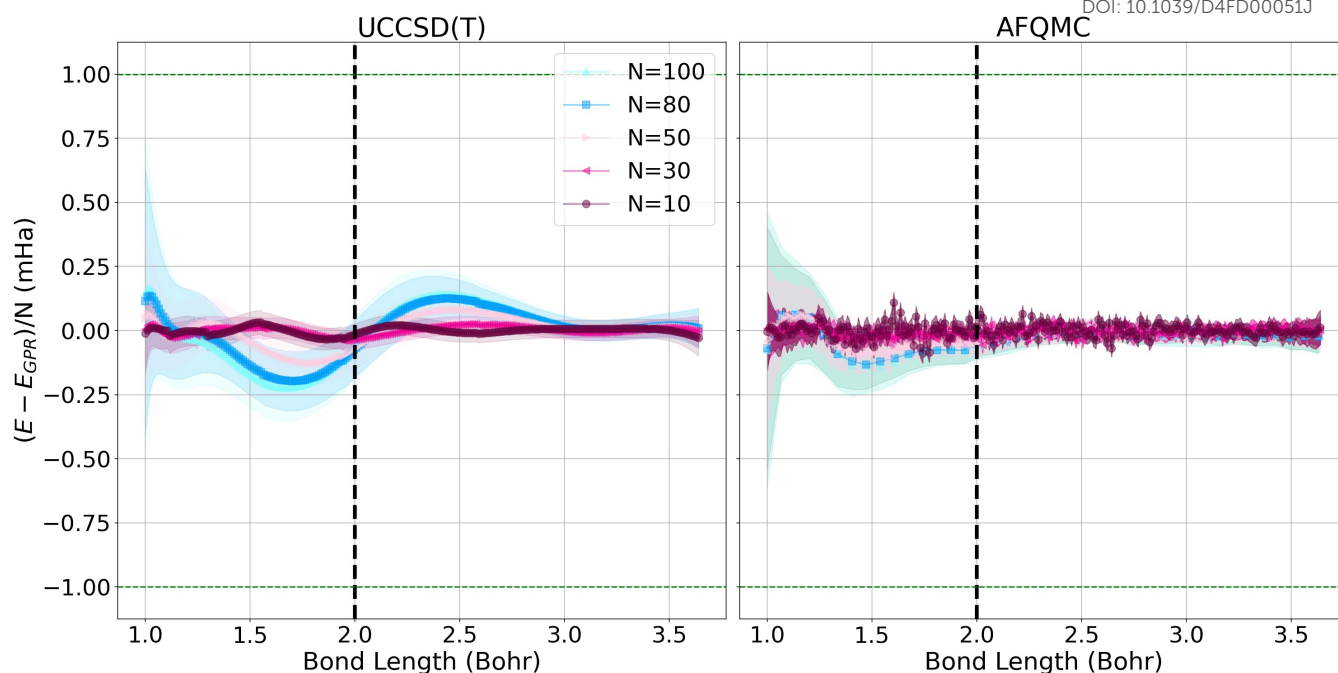


FIG. 3. Energy differences between the calculated UCCSD(T) (left) and AFQMC (right) energies, and their respective  $E_{GPR}$  predictions per atom for hydrogen chains of different lengths in mHa. The green dashed lines depict the bounds of 1 mHa energy differences. The shadows delineate 95% confidence intervals based on the predicted variance. The vertical dashed line denotes the bond length at which the metal-insulator transition occurs.

6

however, noisier than the UCCSD(T) differences, which reflects the stochastic character of AFQMC. The AFQMC-GPR differences are, in general, smaller than the UCCSD(T)-GPR predictions, especially at intermediate bond lengths. Overall, the AFQMC predictions are slightly more accurate and homogeneous at all of the bond lengths studied, likely due to a larger consistency within the AFQMC data.

#### 4. Extrapolation of Chain Energies to the Thermodynamic Limit

Given the sub-milliHartree accuracy of these predictions, we now turn to analyzing the performance of our GPR predictions for extrapolating the energies of very long, yet finite chains that approach the thermodynamic limit. In previous studies,<sup>1</sup> thermodynamic limit predictions were made by assuming the chain energies varied polynomially with  $N^{-1}$ , with orders ranging from 1 to 3 depending upon the convergence speed exhibited by the data.<sup>1</sup> To make use of such scaling laws, a polynomial must be fit to a large enough number of different chain sizes to capture the correct scaling behavior. To compare the performance of our GP regressions against this more conventional fitting procedure, we fit the energies of chains containing 10, 30, and 50 atoms, as was done in Reference 1. We contrasted the extrapolations produced by this polynomial fit with GPR results trained *once* across different bond lengths on chains of 10, 20, and 30 atoms. Indeed, the primary advantage of our method is that we can automat-

ically predict the energy per atom of any chain by computing its global descriptor vector and using the posterior to predict its energy. As an added benefit, the confidence intervals based on the predicted variance provide an estimate of the uncertainty of the prediction, which is not available from typical polynomial regressions.

Figure 4 displays the convergence of the energy per atom to the thermodynamic limit for four representative bond lengths. The circles denote the UCCSD(T) calculations while the squares represent the GPR predictions at each size. As before, the shadows delineate 95% confidence intervals on the GPR calculations. The green dashed line denotes the polynomial regression at the given bond length and the red triangle represents the energy in the thermodynamic limit taken from Reference 1. The GPR prediction of the energy in the thermodynamic limit is made using a chain of 5000 hydrogen atoms. As an illustration of the speed of our regression technique, producing the descriptors for the 5000-atom chain took about 2 minutes on an Intel Core i7-8550U (Turbo 4.0 GHz, 4 Cores, 8 Threads) laptop. We note that the differences between the thermodynamic limit predictions made by the reference regression<sup>1</sup> and the polynomial regression performed on our dataset simply reflect the small differences between the two different databases. The GPR predictions are in good agreement with the reference and polynomial regressions, deviating most for bond lengths near the equilibrium bond length (around 1.8 bohr) where the convergence is less linear. Note that the energies converge one to two orders of magnitude more rapidly



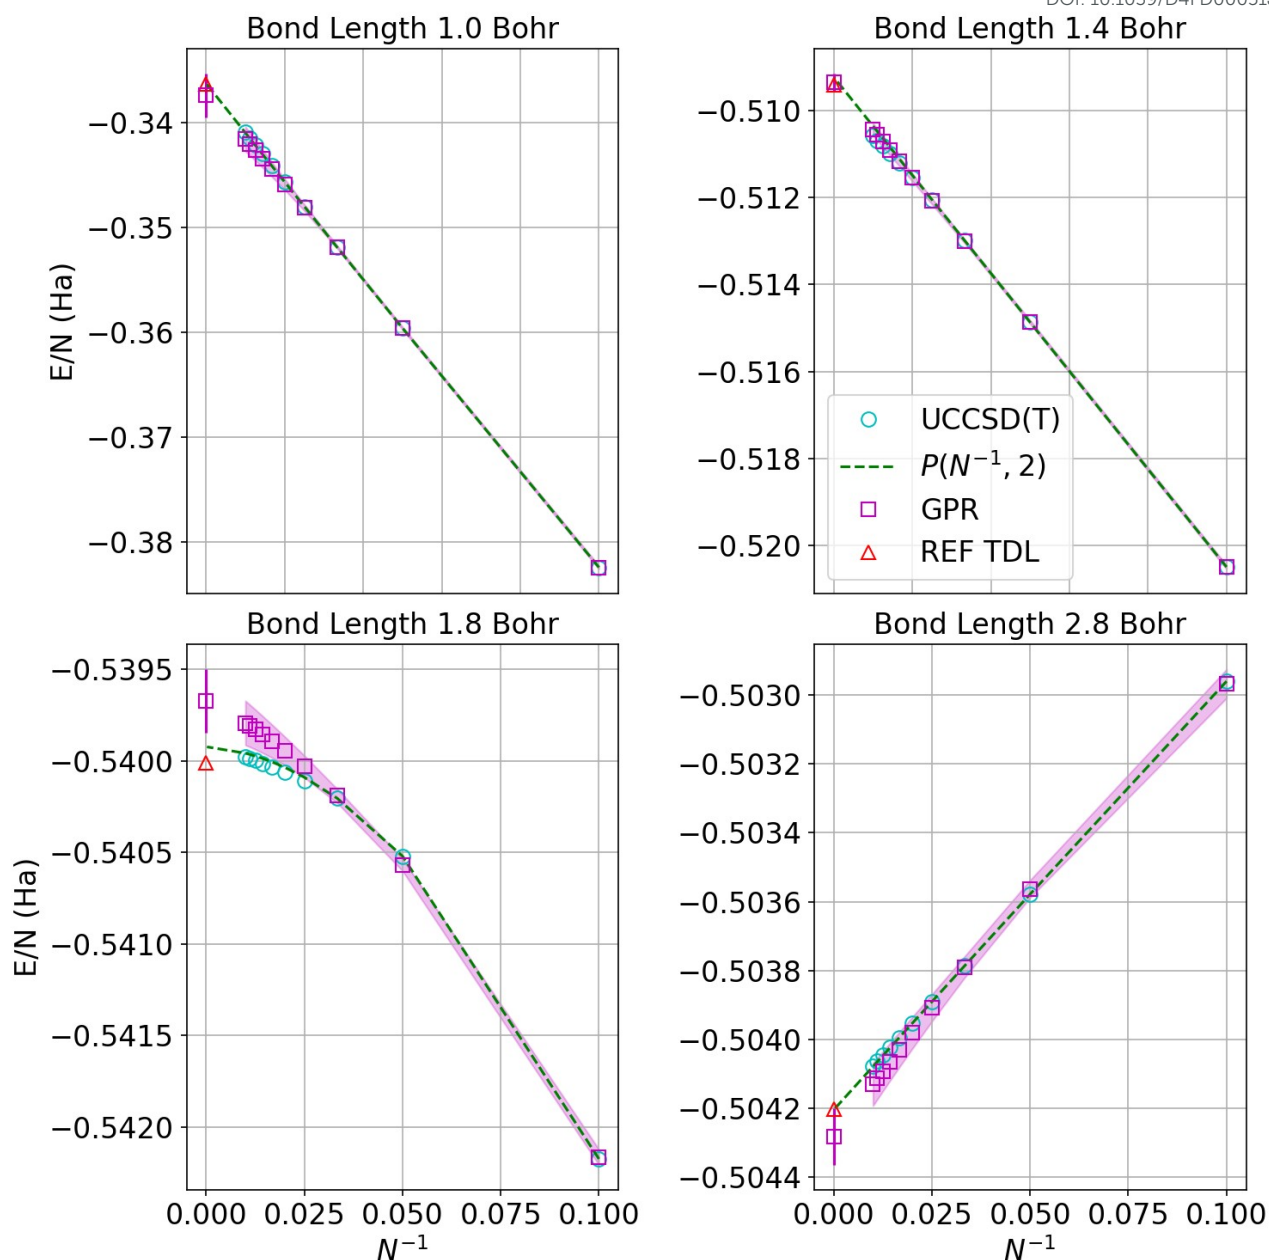


FIG. 4. Predictions of the energy per atom in the thermodynamic limit vs.  $N^{-1}$  based on UCCSD(T) results for hydrogen chains with bond lengths of 1.0, 1.4, 1.8, and 2.8 Bohr. Cyan circles denote UCCSD(T) calculations, while maroon squares denote the GPR predictions. The shadow depicts 95% confidence intervals, and the dashed lines depict the polynomial regression of second order at each bond length. The triangle represents the energy in the thermodynamic limit computed in Reference 1. The gap between points at small values of  $N^{-1}$  corresponds to chains between 100 and 5000 atoms, which are prohibitive to model even using less expensive theories.

at larger bond lengths because the long-range Coulomb interaction is weaker at larger bond lengths, as described earlier.

Figure 5 similarly exhibits the convergence to the thermodynamic limit for the AFQMC database and its respective GPR predictions. One of the most noticeable differences relative to the UCCSD(T) calculations is that the AFQMC predictions seem more linear close to the equilibrium bond

length. This means that the AFQMC calculations can more accurately resolve small, sub-milliHartree differences in the energies as a function of system size and therefore so can the AFQMC-based GPR.

The left panel of Figure 6 presents the energy of the hydrogen chains as a function of bond length directly calculated using UHF and UCCSD(T) for 100-atom chains, as well



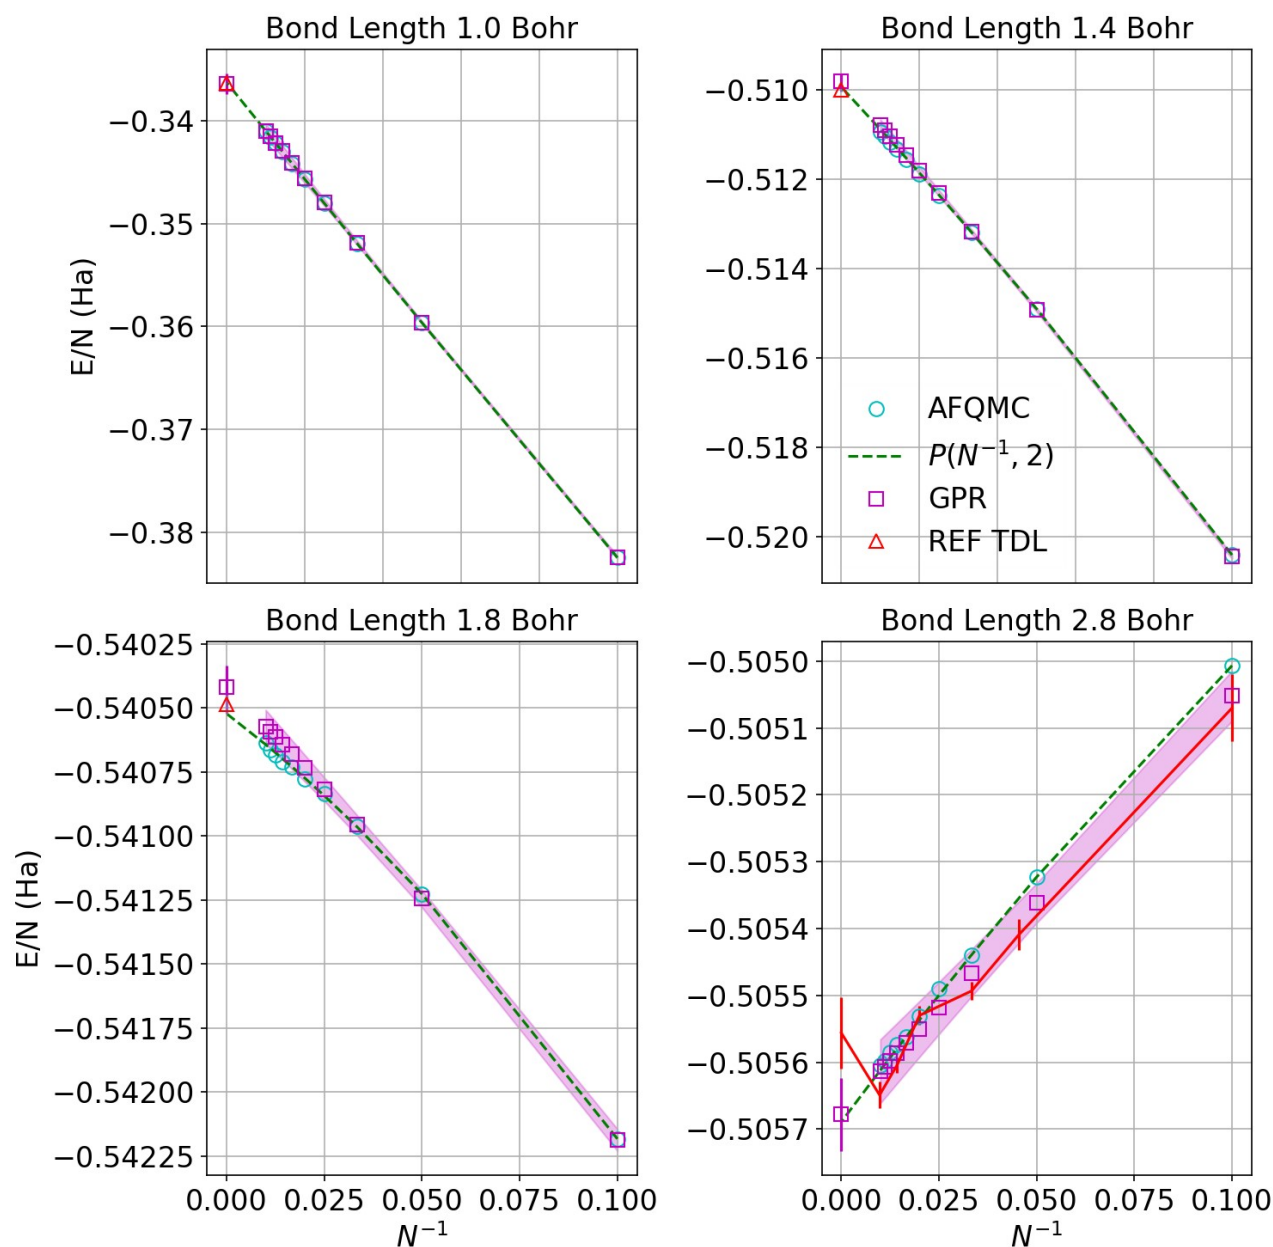


FIG. 5. Predictions of the energy per atom in the thermodynamic limit vs.  $N^{-1}$  based on AFQMC results for hydrogen chains with bond lengths of 1.0, 1.4, 1.8, and 2.8 Bohr. Cyan circles denote direct AFQMC calculations, while the maroon squares denote the GPR predictions. The shadow depicts 95% confidence intervals, and the dashed lines depict the polynomial regression of second order at each bond length. The triangle represents the energy in the thermodynamic limit computed in Reference 1. The gap between points at small values of  $N^{-1}$  corresponds to chains between 100 and 5000 atoms, which are too prohibitive to compute using even less expensive theories. Note that the TDL of Reference 1 (REF TDL) for the bond length of 2.8 Bohr was replaced by our TDL extrapolation using a polynomial regression because the reference value seemed to be in disagreement with the rest of the reference's data at that bond length.

as the reference<sup>1</sup> and GPR predictions in the thermodynamic limit. The energy differences between the  $N = 100$  UCCSD(T), reference, and GPR predictions are hardly perceptible. The right panel of Figure 6 likewise presents the energy as a function of bond length for the largest,  $N = 100$ -atom AFQMC calculations we were able to perform, in addition to the reference and GPR thermodynamic limit

predictions. As in the UCCSD(T) case, the discrepancies are too small to discern at this scale.

To more closely examine how the GPR predictions converge with the number of atoms in the chains, in Figure 7, we plot the difference between the thermodynamic limit predictions of Ref. 1 and our UCCSD(T) (left) and AFQMC-based



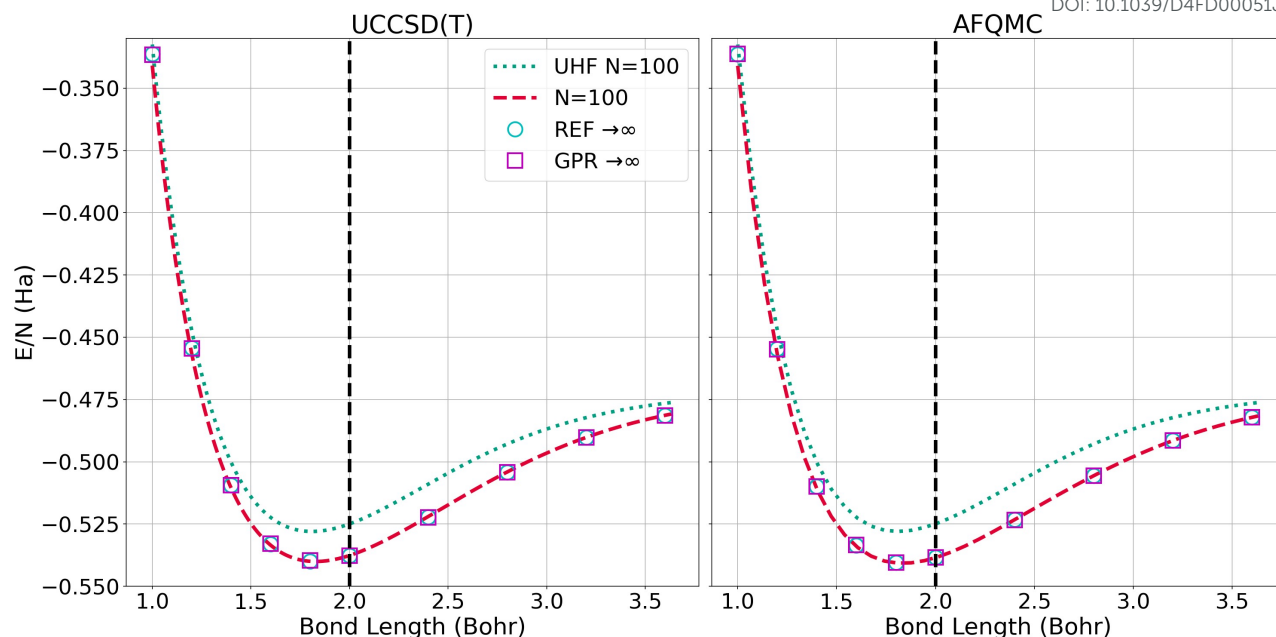


FIG. 6. Energy per atom computed for  $N = 100$  chains and predicted for  $N = \infty$  chains using the UCCSD(T) (left) and AFQMC (right) methods. The “REF  $\rightarrow \infty$ ” is the TDL extrapolation taken from Ref. 1. We plot this reference’s extrapolation so that it can be contrasted with our GPR’s prediction using 5000 atoms.

(right) GPR predictions on the milliHartree scale. For both methods, we take  $N = 5000$  hydrogen chain GPR predictions to be representative of the thermodynamic limit. In the left-hand panel, we also plot UCCSD(T) results for  $N = 200$  hydrogen chains, the largest we could directly simulate, to contrast  $N = 200$  with  $N \rightarrow \infty$  results. We see that, at smaller bond lengths, discrepancies still remain between the  $N = 200$  and  $N \rightarrow \infty$  results, signifying that finite size effects still influence the energies of even  $N = 200$ -length chains.

These discrepancies are also manifested in the larger confidence intervals that accompany the GPR predictions. Even so, GPR predictions at all bond lengths studied possess sub-milliHartree accuracy, and the discrepancies between the different chain length predictions disappear at the longest bond lengths studied. In contrast, the right panel of Figure 7 demonstrates that the AFQMC-based GPR predictions are in much better agreement with Ref. 1’s thermodynamic limit predictions, even at shorter bond lengths. This is in line with the results presented earlier in Figure 3.

Since our calculations were performed with open boundary conditions (OBC), it is also worthwhile to compare our predictions to those produced using the “subtraction trick,”<sup>33</sup> in which the energies of systems of different sizes are subtracted to eliminate surface effects from bulk energies. Figure 8 presents the differences in energy between our GPR predictions of the energies in the thermodynamic limit and those produced using the subtraction trick based on chains of different lengths. The differences in the energies predicted by these approaches is sub-milliHartree at all bond lengths stud-

ied, further demonstrating that our GPR predictions are highly accurate relative to a widely-employed benchmark, while also illustrating the surprising accuracy of the subtraction trick. As the subtraction trick eliminates edge effects from energy predictions, this comparison especially highlights the GPR method’s ability to correct for edge effects. It is satisfying to see that the energies predicted by the subtraction trick performed on chains of lengths 30 and 50, which should yield the most accurate predictions of the subtraction trick calculations, are in the greatest agreement with our GPR predictions, especially at intermediate bond lengths. As before, we see that our GPR predictions are in the greatest agreement with the subtraction trick results at longer bond lengths. Indeed, our GPR predictions almost perfectly agree with all three of the subtraction trick predictions at the longest bond lengths studied. Moreover, our AFQMC-based GPR predictions again converge more rapidly and reliably to the thermodynamic limit with increasing bond length. Overall, Figures 7 and 8 possess very similar features: the GPR predictions overestimate the energies at the shortest bond lengths and then oscillate between under- and overestimating the energies at intermediate bond lengths before coming to agreement at the longest bond lengths. This points to the overwhelming agreement between the polynomial regression and subtraction trick energies. These comparisons also demonstrate that the GPR predictions are not uniformly biased toward over- or underestimating energies.

A more quantitative comparison of the predictions generated by all of these methods can be found in the Supplementary Materials.





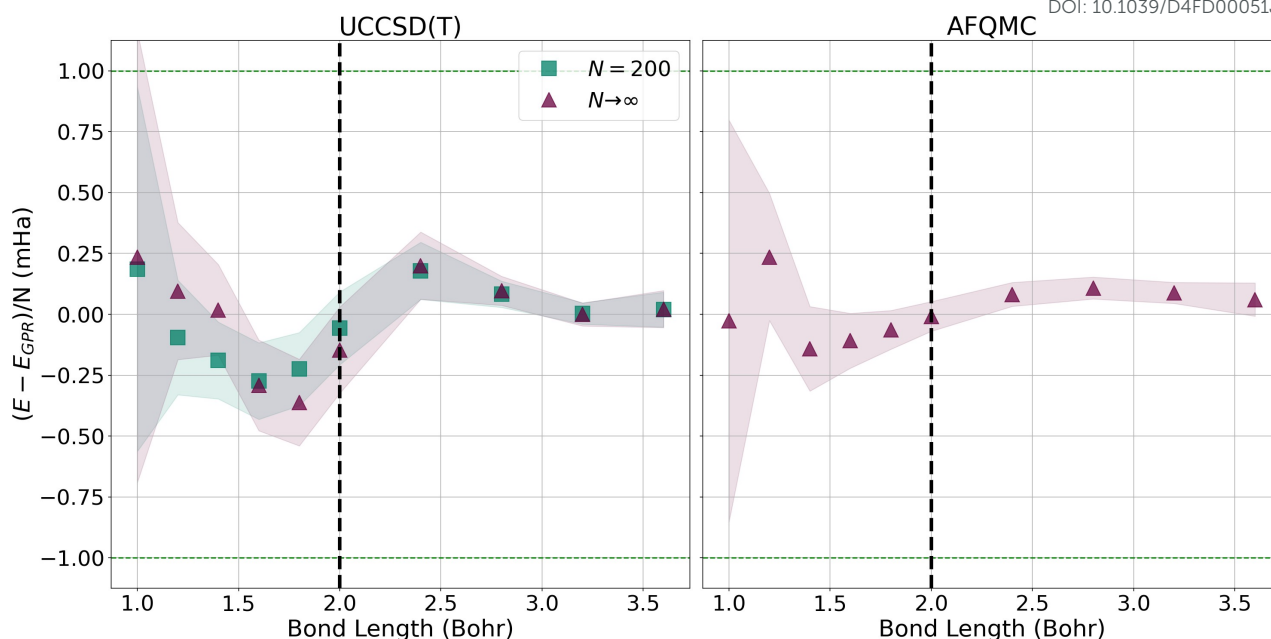


FIG. 7. Differences between the energies predicted by Ref. 1 and the UCCSD(T) (left) and AFQMC (right) GPR-predicted energies in the thermodynamic limit (red triangles). For both the UCCSD(T) and AFQMC plots, we assume that the GPR prediction using 5000 atoms is representative of the GPR prediction in the thermodynamic limit. On the left, we also plot the UCCSD(T) energies for  $N = 200$  hydrogen chains, the largest we could directly simulate. The shadows depict 95% confidence intervals for the GPR predictions.

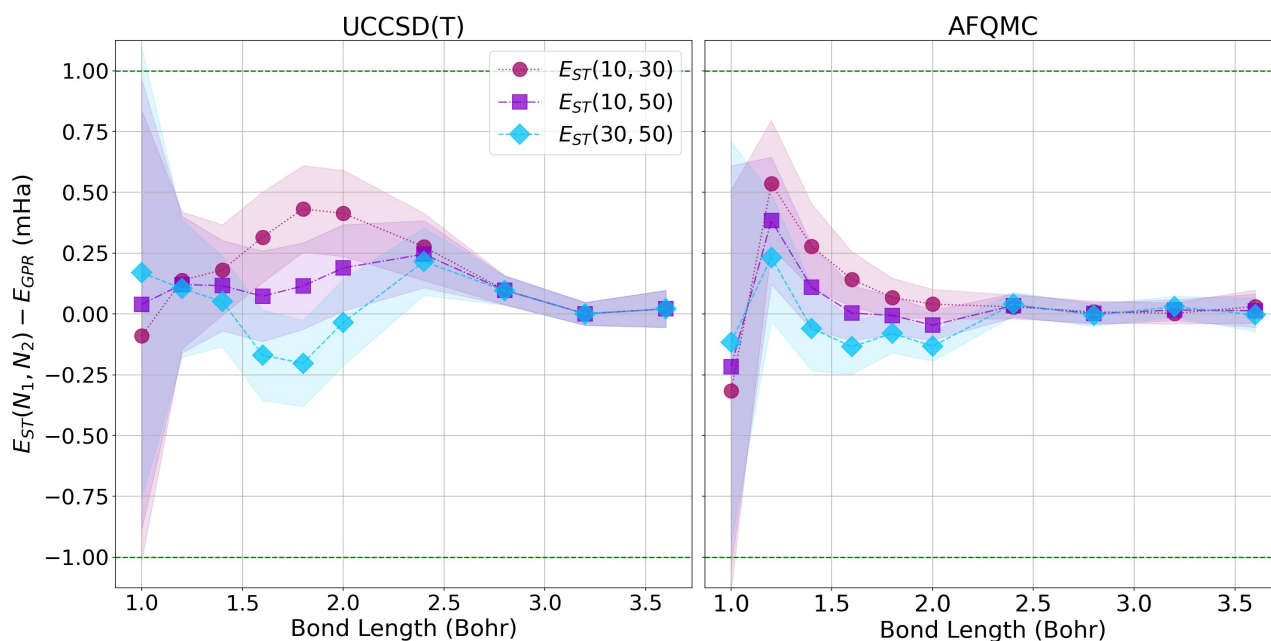


FIG. 8. The difference in energies between our GPR predictions in the thermodynamic limit,  $E_{GPR}(N \rightarrow \infty)$ , and extrapolated energies obtained using the subtraction trick,  $E_{ST}$ . (Left) Differences based upon UCCSD(T) energies; (Right) differences based upon AFQMC energies.

## B. Two-Dimensional, Inhomogeneous Hydrogen Chains

Given the success of GPR at predicting the energies of homogeneously-stretched hydrogen chains in the thermodynamic limit, we next examine the capacity for the same GPR

techniques to predict the energies of inherently heterogeneous chains of hydrogen dimers. As depicted in Figure 9, these chains of hydrogen dimers are described by two key distances: the intra-dimer bond distance,  $a$ , and the inter-dimer bond distance,  $b$ . In the following, we generally fix the intradimer dis-



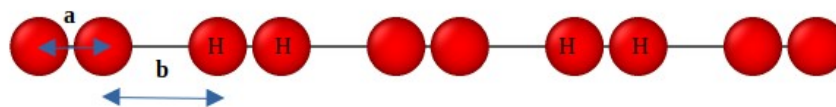


FIG. 9. Illustration of the linear chains of hydrogen dimers studied in this work with intradimer distance,  $a$ , and interdimer distance,  $b$ .

tance,  $a$ , between 1.0 and 3.5 bohr, and vary the interdimer distance between 1.0 and  $a$  bohr, maintaining open boundary conditions. While these chains of dimers enable us to retain the same periodicity present in our earlier homogeneous chains, they also enable us to purposefully and controllably introduce heterogeneity into our systems that complicates our prediction problem. Indeed, these chains of dimers manifest several levels of correlation when stretched, typically necessitating the use of advanced quantum chemistry methods to make high-accuracy energy predictions.<sup>91</sup>

To study the performance of our GPR algorithm on these chains, we generate a database of dimer chain energies starting from UHF calculations with single Slater determinants that we again input into either CCSD(T) or AFQMC calculations. We model our hydrogen atoms using the minimal STO-6G basis set given the steep computational cost of the system with increasing system size. Chains of 5, 10, and 15 dimers for a total of 176 configurations were employed for training. The remaining 315 configurations of chains consisting of 20 to 50 dimers were subsequently used for testing and validation. The same atomic environment descriptors previously employed for the homogeneous chains were also employed here.

The energy surfaces for chains consisting of  $N = 30$ , 50, and 100 atoms are depicted in Figure 10. It can be seen that the GPR predictions are in qualitative agreement with the AFQMC database values, both for short chains ( $N=10$ ) and long chains approaching the thermodynamic limit ( $N=100$ ). In particular, GPR is able to well describe both the energy minimum around  $a=1.5$  bohr,  $b=3.25$  bohr, and highly stretched chains with both  $a$  and  $b$  greater than 3 bohr. More detailed slices of the potential energy surface for several values of  $a$  are depicted in Figure 11.

As is apparent from these plots, the approach to large interdimer separations is highly dependent upon the intra-dimer separation: for small  $a$ , the approach is steeper than for large  $a$ . This behavior is a sign of correlation between the  $a$  and  $b$  values and is non-trivial, given the seeming simplicity of the model. This makes the model a useful testbed for multidimensional extrapolations, as further discussed in the Supplementary Materials.

To visualize the energy surface, as shown in the Figure 12, we use triangulation over the sample points and then Delaunay smoothing. The thermodynamic limit was estimated using GPR regression on  $N=5000$  atoms. On the left, we present the 3D energy surface for a chain of 15 dimers; the black dots denote the energies estimated by GPR in the thermodynamic limit. On the right, we provide a heat map corresponding to the plot on the left annotated with iso-energy contour lines

predicted using GPR for systems of different sizes. In particular, the red and yellow dashed lines denote the energies for chains comprised of 15 and 50 dimers, respectively. The errors on these energies are all less than 1 mHa, which is within chemical accuracy.

This plot underscores how the contours change or shift with system size. We can see that the largest differences between the contours occur near the minimum of the plot around an intra-dimer distance of 1.5 bohr and an interdimer distance of 3 bohr. In this region of the surface, the  $N = 30$  contours differ significantly from the  $N = 100$  contours, which nearly align with the thermodynamic limit contours, suggesting that 100 atoms are nearly enough to converge simulations of this system to their TDL. Away from this minimum, the contours for all three system sizes concur, demonstrating that the system experiences weaker finite size effects for these parameters. GPR's success extrapolating the energies of this non-trivial, multidimensional model suggest that it is likely to have similar success on the more complex models and solids of interest to the wider scientific community.

## V. DISCUSSION OF RESULTS

Although we employed Gaussian Process Regression in this work, a wide range of other machine learning approaches, including artificial neural networks, could also be used to perform these extrapolations. We opted to employ kernel methods like Gaussian processes because they are non-parametric and make use of Bayesian inference at a comparatively low,  $O(N_t^3)$  cost, where  $N_t$  is the size of the training set.<sup>17</sup> It has been proposed<sup>92</sup> as a rule of thumb to use  $N_t = 10 \times d$ , where  $d$  is the dimension of the feature space, to train a GPR. In contrast, neural network-based approaches involve matrix-vector multiplications that scale with the number of neurons in the network,  $N_n$ , and the dimension of the input vector,  $d$ , as  $O(N_n d)$ . If the number of neurons in the network is small, this implies that neural networks are less expensive to employ than GPR. However, neural networks typically necessitate the use of non-linear activation functions that may increase their overall cost. More importantly, neural networks often suffer from overfitting if care is not taken to reoptimize their number of nodes or structures. Overfitting is much less of a concern for GPR since GPR with the same kernel but more training points is guaranteed to be more accurate. In practice, NNs use at least 2 to 3 orders of magnitude more training data than GPR.<sup>17,32</sup> When training data is scarce - as it usually is when many-body electronic structure calculations are involved - GPR-based techniques hence become the method of choice.<sup>32</sup> One






FIG. 10. Energy surfaces,  $E/N$ , predicted for chains consisting of (left) 30, (center) 50, and (right) 100 atoms (15, 25, and 50 dimers, respectively) for several  $a$  and  $b$  values. AFQMC energies are given by the cyan circles, while GPR predictions are given by the maroon triangles.




FIG. 11. Comparison of GPR and AFQMC predictions for different intra-dimer bond lengths as a function of inter-dimer distances. The  $N=10$ , 50, and 100-atom data are all provided by AFQMC.

may also ask whether using GPR on these low-dimensional data sets is more sophisticated than necessary and if other, less sophisticated regression techniques based on a small number of parameters could instead be employed. As demonstrated in the Supplementary Material, we have compared the performance of our GPR approach to that of Bayesian Multivariate Adaptive Regression Splines, a spline-based technique, and found that our GPR approach can extrapolate with significantly greater accuracy. We moreover show that, while one can extrapolate using a few simple parameters, this extrapolation is not readily generalizable to more complex situations in which the parameters to use are less obvious. Lastly, as illustrated throughout this manuscript, GPR inherently quantifies uncertainties, which are critical for being able to determine its accuracy relative to that of other methods.

Our studies of low-dimensional hydrogen chains naturally beg the question of how well our techniques can be generalized to more realistic multidimensional solids that are accompanied by an even more rapid growth in computational

expense. Much like other GAP methods, our approach should be readily generalizable to higher dimensional systems, given sufficient data and high-quality features. Indeed, here, we took the first step toward demonstrating this by applying our model to both a one-dimensional and a nontrivial two-dimensional system, and in a previous preprint, we demonstrated how a similar GPR-based approach could be leveraged to predict the energies of 3D alloys.<sup>93</sup> The key challenge associated with higher-dimensional predictions is the curse of dimensionality: the higher the dimensionality of the space, the more data that is needed for training to learn the larger space with sufficient accuracy to make effective comparisons between different atomic environments. The resulting increase in cost can be slowed through a more judicious selection of features and design of kernels. CUR<sup>88</sup> decompositions and Kernelized Principal Covariates Regression<sup>90</sup> are excellent alternatives for identifying the most relevant features, which can significantly reduce the dimension of the descriptors of a given data set. More effective kernels may also



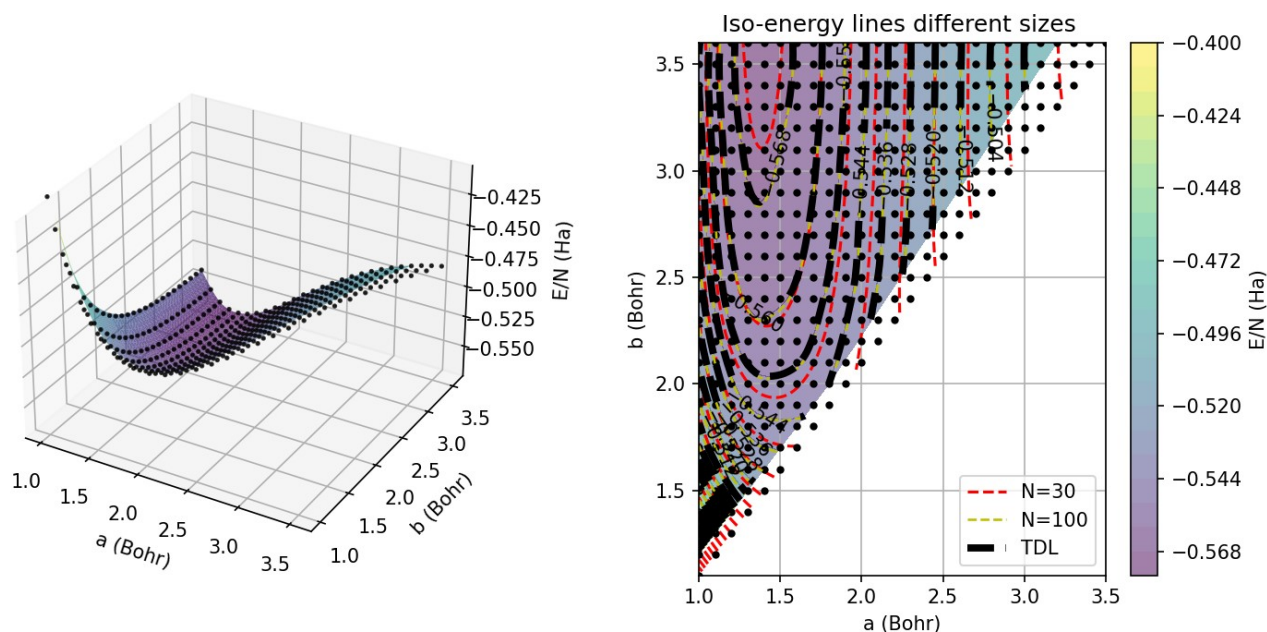


FIG. 12. (Left) Interpolated GPR potential energy surface,  $E/N$ , for a chain of 15 dimers as a function of  $a$  and  $b$ . The black dots denote the AFQMC training points from the database. (Right) Iso-contours of the energy surface,  $E/N$ , interpolated for chains of  $N = 30$  (red) and  $N = 100$  (yellow) atoms, and in the thermodynamic limit (dashed black lines). The color map denotes the GPR predictions of the energy in the thermodynamic limit.

be constructed through approaches that recursively evaluate the differences between structures.<sup>69</sup> For example, De *et al.* proposed kernels based on regularized structure matching to optimize the comparisons between the atomic environments of different structures.<sup>69</sup> Thus, with further technical developments, we believe that the techniques presented here should be readily generalizable to the even larger, more complicated solids that they would most benefit.

## VI. CONCLUSIONS

In summary, in this work, we have presented a Gaussian Process Regression-based approach for predicting the many-body energies of hydrogen chains, the simplest examples of *ab initio* solids, in the thermodynamic limit. We have shown that, by training on databases of the energies of short (10-30-atom) homogeneous and inhomogeneous hydrogen chains with varying intra- and inter-dimer distances, we can predict the energies of these chains in the thermodynamic limit with sub-milliHartree accuracy relative to predictions made by alternative extrapolation techniques. These alternative techniques, including polynomial regressions and the “subtraction trick,” typically necessitate computing the energies of chains much longer than the chains employed in our training sets. As such, our approach enables the highly accurate prediction of the energies of solids in the thermodynamic limit based upon relatively small systems, and hence, much less expensive calculations. Unlike many finite size extrapolation techniques which apply to systems with only certain geometries,

densities, and/or dimensionality, as demonstrated by the easy generalizability of our method to both homogeneous and inhomogeneous chains, our approach is largely agnostic to the physical characteristics of the system studied; as long as there is sufficient and representative training data, our approach can be applied, making it particularly useful for some of the more complex systems of modern interest, such as those at interfaces or having irregular geometries.

## VII. DATA AVAILABILITY

The data that support the findings of this study are openly available online at [https://github.com/josuelandinez/LHC\\_Database](https://github.com/josuelandinez/LHC_Database).

## VIII. SUPPLEMENTARY MATERIAL

A supplementary materials document has been provided that contains comparisons of the accuracy of our GPR methods to those of other regression techniques and tables of the energies predicted by the different approaches described here that underlie many of the plots provided in the main text.

## IX. ACKNOWLEDGEMENTS

E.J.L.B. would like to thank Amit Samanta for insightful discussions and Qiming Sun for his kind assistance with per-





ods. *Condensed Matter Physics*, 26(3):33701, 2023. ISSN 1607-324X. doi:10.5488/cmp.26.33701. URL <http://dx.doi.org/10.5488/CMP.26.33701>.

forming calculations in PySCF. E.J.L.B. (concept, modeling, analysis, manuscript preparation), A.L. (manuscript preparation), and B.R. (concept, mentoring, manuscript preparation) graciously acknowledge support from U.S. Department of Energy, Office of Science, Basic Energy Sciences Award #DE-FOA-0001912. B.R. (concept) also received support from U.S. National Science Foundation grant OIA-1921199 and the Research Corporation of America (concept). This research was conducted using computational resources and services at the Center for Computation and Visualization, Brown University.

## X. REFERENCES

- Mario Motta, David M. Ceperley, Garnet Kin-Lic Chan, John A. Gomez, Emanuel Gull, Sheng Guo, Carlos A. Jiménez-Hoyos, Tran Nguyen Lan, Jia Li, Fengjie Ma, Andrew J. Millis, Nikolay V. Prokof'ev, Ushnish Ray, Gustavo E. Scuseria, Sandro Sorella, Edwin M. Stoudenmire, Qiming Sun, Igor S. Tupitsyn, Steven R. White, Dominika Zgid, and Shiwei Zhang. *Phys. Rev. X*, 7:031059, Sep 2017. doi:10.1103/PhysRevX.7.031059. URL <https://link.aps.org/doi/10.1103/PhysRevX.7.031059>.
- G. Kresse and J. Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54:11169–11186, Oct 1996. doi:10.1103/PhysRevB.54.11169. URL <https://link.aps.org/doi/10.1103/PhysRevB.54.11169>.
- Stefan Goedecker. Linear scaling electronic structure methods. *Rev. Mod. Phys.*, 71:1085–1123, Jul 1999. doi:10.1103/RevModPhys.71.1085. URL <https://link.aps.org/doi/10.1103/RevModPhys.71.1085>.
- Stephan Mohr, Laura E. Ratcliff, Luigi Genovese, Damien Caliste, Paul Boulanger, Stefan Goedecker, and Thierry Deutsch. Accurate and efficient linear scaling dft calculations with universal applicability. *Phys. Chem. Chem. Phys.*, 17:31360–31370, 2015. doi:10.1039/C5CP00437C. URL <http://dx.doi.org/10.1039/C5CP00437C>.
- Markus Holzmann, Raymond C. Clay, Miguel A. Morales, Norm M. Tubman, David M. Ceperley, and Carlo Pierleoni. Theory of finite size effects for electronic quantum monte carlo calculations of liquids and solids. *Phys. Rev. B*, 94:035126, Jul 2016. doi:10.1103/PhysRevB.94.035126. URL <https://link.aps.org/doi/10.1103/PhysRevB.94.035126>.
- N. D. Drummond, R. J. Needs, A. Sorouri, and W. M. C. Foulkes. Finite-size errors in continuum quantum monte carlo calculations. *Phys. Rev. B*, 78:125106, Sep 2008. doi:10.1103/PhysRevB.78.125106. URL <https://link.aps.org/doi/10.1103/PhysRevB.78.125106>.
- Sam Azadi and W. M. C. Foulkes. Systematic study of finite-size effects in quantum monte carlo calculations of real metallic systems. *The Journal of Chemical Physics*, 143(10):102807, 2015.
- Tina N. Mihm, Bingdi Yang, and James J. Shepherd. Power laws used to extrapolate the coupled cluster correlation energy to the thermodynamic limit. *Journal of Chemical Theory and Computation*, 17(5):2752–2758, 2021.
- A. Baldereschi. Mean-value point in the brillouin zone. *Phys. Rev. B*, 7:5212–5215, Jun 1973. doi:10.1103/PhysRevB.7.5212. URL <https://link.aps.org/doi/10.1103/PhysRevB.7.5212>.
- C. Lin, F. H. Zong, and D. M. Ceperley. Twist-averaged boundary conditions in continuum quantum monte carlo algorithms. *Phys. Rev. E*, 64:016702, Jun 2001. doi:10.1103/PhysRevE.64.016702. URL <https://link.aps.org/doi/10.1103/PhysRevE.64.016702>.
- W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum monte carlo simulations of solids. *Rev. Mod. Phys.*, 73:33–83, Jan 2001. doi:10.1103/RevModPhys.73.33. URL <https://link.aps.org/doi/10.1103/RevModPhys.73.33>.
- Hendra Kwee, Shiwei Zhang, and Henry Krakauer. Finite-size correction in many-body electronic structure calculations. *Phys. Rev. Lett.*, 100:126404, Mar 2008. doi:10.1103/PhysRevLett.100.126404. URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.126404>.
- V. Gorelov, Y. Yang, M. Ruggeri, D. M. Ceperley, C. Pierleoni, and M. Holzmann. Neutral band gap of carbon by quantum monte carlo meth-
- Vitaly Gorelov, Markus Holzmann, David M. Ceperley, and Carlo Pierleoni. Electronic excitation spectra of molecular hydrogen in phase i from quantum monte carlo and many-body perturbation methods, 2023.
- M. Colomé-Tatché, S. I. Matveenko, and G. V. Shlyapnikov. Finite-size effects for the gap in the excitation spectrum of the one-dimensional hubbard model. *Phys. Rev. A*, 81:013611, Jan 2010. doi:10.1103/PhysRevA.81.013611. URL <https://link.aps.org/doi/10.1103/PhysRevA.81.013611>.
- Ke Liao and Andreas Grüneis. Communication: Finite size correction in periodic coupled cluster theory calculations of solids. *The Journal of Chemical Physics*, 145(14):141102, 2016.
- Pankaj Mehta, Marin Bukov, Ching-Hao Wang, Alexandre G.R. Day, Clint Richardson, Charles K. Fisher, and David J. Schwab. A high-bias, low-variance introduction to machine learning for physicists. *Physics Reports*, 810:1 – 124, 2019. ISSN 0370-1573. doi: <https://doi.org/10.1016/j.physrep.2019.03.001>. URL <http://www.sciencedirect.com/science/article/pii/S0370157319300766>. A high-bias, low-variance introduction to Machine Learning for physicists.
- Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91:045002, Dec 2019. doi:10.1103/RevModPhys.91.045002. URL <https://link.aps.org/doi/10.1103/RevModPhys.91.045002>.
- Chandramouli Nyshadham, Matthias Rupp, Brayden Bekker, Alexander V. Shapeev, Tim Mueller, Conrad W. Rosenbrock, Gábor Csányi, David W. Wingate, and Gus L. W. Hart. Machine-learned multi-system surrogate models for materials prediction. *npj Computational Materials*, 5(1):51, Apr 2019. ISSN 2057-3960. doi:10.1038/s41524-019-0189-9. URL <https://doi.org/10.1038/s41524-019-0189-9>.
- Jonathan Schmidt, Mário R. G. Marques, Silvana Botti, and Miguel A. L. Marques. Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5(1):83, Aug 2019. ISSN 2057-3960. doi:10.1038/s41524-019-0221-0. URL <https://doi.org/10.1038/s41524-019-0221-0>.
- Alexandru B. Georgescu, Peiwen Ren, Aubrey R. Tolland, Shengtong Zhang, Kyle D. Miller, Daniel W. Apley, Elsa A. Olivetti, Nicholas Wagner, and James M. Rondinelli. Database, features, and machine learning model to identify thermally driven metal–insulator transition compounds. *Chemistry of Materials*, 33(14):5591–5605, 2021. doi:10.1021/acs.chemmater.1c00905. URL <https://doi.org/10.1021/acs.chemmater.1c00905>.
- Juan Carrasquilla. Machine learning for quantum matter. *Advances in Physics: X*, 5(1):1797528, 2020. doi:10.1080/23746149.2020.1797528. URL <https://doi.org/10.1080/23746149.2020.1797528>.
- David Pfau, James S. Spencer, Alexander G. D. G. Matthews, and W. M. C. Foulkes. Ab initio solution of the many-electron schrödinger equation with deep neural networks. *Phys. Rev. Res.*, 2:033429, Sep 2020. doi:10.1103/PhysRevResearch.2.033429. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.2.033429>.
- Jan Hermann, Zeno Schätzle, and Frank Noé. Deep-neural-network solution of the electronic schrödinger equation. *Nature Chemistry*, 12(10):891–897, Oct 2020. ISSN 1755-4349. doi:10.1038/s41557-020-0544-y. URL <https://doi.org/10.1038/s41557-020-0544-y>.
- Aldo Glielmo, Yannic Rath, Gábor Csányi, Alessandro De Vita, and George H. Booth. Gaussian process states: A data-driven representation of quantum many-body physics. *Phys. Rev. X*, 10:041026, Nov 2020. doi:10.1103/PhysRevX.10.041026. URL <https://link.aps.org/doi/10.1103/PhysRevX.10.041026>.
- Yannic Rath, Aldo Glielmo, and George H. Booth. A bayesian inference framework for compression and prediction of quantum states. *The Journal of Chemical Physics*, 153(12):124108, 2020. doi:10.1063/5.0024570. URL <https://doi.org/10.1063/5.0024570>.
- Hongwei Niu, Yubo Yang, Scott Jensen, Markus Holzmann, Carlo Pierleoni, and David M. Ceperley. Stable solid molecular hydrogen above 900 k from a machine-learned potential trained with diffusion quantum monte carlo. *Phys. Rev. Lett.*, 130:076102, Feb 2023. doi:10.1103/PhysRevLett.130.076102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.130.076102>.



- 10.1103/PhysRevLett.130.076102.
- <sup>28</sup>Dian Wu, Riccardo Rossi, Filippo Vicentini, and Giuseppe Carleo. From tensor-network quantum states to tensorial recurrent neural networks. *Phys. Rev. Res.*, 5:L032001, Jul 2023. doi:10.1103/PhysRevResearch.5.L032001. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.5.L032001>.
- <sup>29</sup>Li Li, Thomas E. Baker, Steven R. White, and Kieron Burke. Pure density functional for strong correlation and the thermodynamic limit from machine learning. *Phys. Rev. B*, 94:245129, Dec 2016. doi:10.1103/PhysRevB.94.245129. URL <https://link.aps.org/doi/10.1103/PhysRevB.94.245129>.
- <sup>30</sup>Laura Weiler, Tina N. Mihm, and James J. Shepherd. Machine learning for a finite size correction in periodic coupled cluster theory calculations. *The Journal of Chemical Physics*, 156(20):204109, 05 2022. ISSN 0021-9606. doi:10.1063/5.0086580. URL <https://doi.org/10.1063/5.0086580>.
- <sup>31</sup>C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. Adaptive computation and machine learning series. University Press Group Limited, 2006. ISBN 9780262182539. URL <https://books.google.com.co/books?id=vWtwQgAACAAJ>.
- <sup>32</sup>Aditya Kamath, Rodrigo A. Vargas-Hernández, Roman V. Krems, Tucker Carrington, and Sergei Manzhos. Neural networks vs gaussian process regression for representing potential energy surfaces: A comparative study of fit quality and vibrational spectrum accuracy. *The Journal of Chemical Physics*, 148(24):241702, 2018.
- <sup>33</sup>E.M. Stoudenmire and Steven R. White. Studying two-dimensional systems with the density matrix renormalization group. *Annual Review of Condensed Matter Physics*, 3(1):111–128, 2012. doi:10.1146/annurev-conmatphys-020911-125018. URL <https://doi.org/10.1146/annurev-conmatphys-020911-125018>.
- <sup>34</sup>Peter J. Reynolds, Jan Tobochnik, and Harvey Gould. Diffusion Quantum Monte Carlo. *Computer in Physics*, 4(6):662–668, 11 1990. ISSN 0894-1866. doi:10.1063/1.4822960. URL <https://doi.org/10.1063/1.4822960>.
- <sup>35</sup>Tao Pang. Diffusion Monte Carlo: A powerful tool for studying quantum many-body systems. *American Journal of Physics*, 82(10):980–988, 10 2014. ISSN 0002-9505. doi:10.1119/1.4890824. URL <https://doi.org/10.1119/1.4890824>.
- <sup>36</sup>J. C. Greer. Estimating full configuration interaction limits from a Monte Carlo selection of the expansion space. *The Journal of Chemical Physics*, 103(5):1821–1828, 08 1995. doi:10.1063/1.469756. URL <https://doi.org/10.1063/1.469756>.
- <sup>37</sup>Shiwei Zhang and Henry Krakauer. Quantum monte carlo method using phase-free random walks with Slater determinants. *Phys. Rev. Lett.*, 90:136401, Apr 2003. doi:10.1103/PhysRevLett.90.136401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.90.136401>.
- <sup>38</sup>Joonho Lee, Hung Q. Pham, and David R. Reichman. Twenty years of auxiliary-field quantum monte carlo in quantum chemistry: An overview and assessment on main group chemistry and bond-breaking. *Journal of Chemical Theory and Computation*, 18(12):7024–7042, 2022. doi:10.1021/acs.jctc.2c00802. URL <https://doi.org/10.1021/acs.jctc.2c00802>. PMID: 36255074.
- <sup>39</sup>Cancan Huang and Brenda M. Rubenstein. Machine learning diffusion monte carlo forces. *The Journal of Physical Chemistry A*, 127(1):339–355, 2023. doi:10.1021/acs.jpca.2c05904. URL <https://doi.org/10.1021/acs.jpca.2c05904>. PMID: 36576803.
- <sup>40</sup>Juha Tiihonen, III Clay, Raymond C., and Jaron T. Krogel. Toward quantum Monte Carlo forces on heavier ions: Scaling properties. *The Journal of Chemical Physics*, 154(20):204111, 05 2021. ISSN 0021-9606. doi:10.1063/5.0052266. URL <https://doi.org/10.1063/5.0052266>.
- <sup>41</sup>Kevin Ryczko, Jaron T. Krogel, and Isaac Tamblyn. Machine Learning Diffusion Monte Carlo Energies. *Journal of Chemical Theory and Computation*, 18(12):7695–7701, 2022. doi:10.1021/acs.jctc.2c00483. URL <https://doi.org/10.1021/acs.jctc.2c00483>. eprint: <https://doi.org/10.1021/acs.jctc.2c00483>.
- <sup>42</sup>Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007. doi:10.1103/PhysRevLett.98.146401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.98.146401>.
- <sup>43</sup>Weiluo Ren, Weizhong Fu, Xiaojie Wu, and Ji Chen. Towards the ground state of molecules via diffusion monte carlo on neural networks. *Nature Communications*, 14(1):1860, Apr 2023. ISSN 2041-1723. doi:10.1038/s41467-023-37609-3. URL <https://doi.org/10.1038/s41467-023-37609-3>.
- <sup>44</sup>Michael Scherbela, Rafael Reisenhofer, Leon Gerard, Philipp Marquetand, and Philipp Grohs. Solving the electronic schrödinger equation for multiple nuclear geometries with weight-sharing deep neural networks. *Nature Computational Science*, 2(5):331–341, May 2022. doi:10.1038/s43588-022-00228-x. URL <https://doi.org/10.1038/s43588-022-00228-x>.
- <sup>45</sup>Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and Weinan E. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.*, 120:143001, Apr 2018. doi:10.1103/PhysRevLett.120.143001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.143001>.
- <sup>46</sup>Note1. Without corrections, QMC methods typically possess an infinite statistical variance,<sup>94</sup> and even with corrections, can scale quite steeply (as  $Z_{eff}^{6.5}$  or greater, where  $Z_{eff}$  is the effective nuclear charge).
- <sup>47</sup>Bing Huang, O. Anatole von Lilienfeld, Jaron T. Krogel, and Anouar Benali. Toward dmc accuracy across chemical space with scalable -qml. *Journal of Chemical Theory and Computation*, 19(6):1711–1721, 2023. doi:10.1021/acs.jctc.2c01058. URL <https://doi.org/10.1021/acs.jctc.2c01058>. PMID: 36857531.
- <sup>48</sup>Peter Broecker, Juan Carrasquilla, Roger G. Melko, and Simon Trebst. Machine learning quantum phases of matter beyond the fermion sign problem. *Scientific Reports*, 7(1):8823, Aug 2017. ISSN 2045-2322. doi:10.1038/s41598-017-09098-0. URL <https://doi.org/10.1038/s41598-017-09098-0>.
- <sup>49</sup>David M. Ceperley, Scott Jensen, Yubo Yang, Hongwei Niu, Carlo Pierleoni, and Markus Holzmann. Training models using forces computed by stochastic electronic structure methods. *arXiv e-prints*, art. arXiv:2310.15994, October 2023. doi:10.48550/arXiv.2310.15994.
- <sup>50</sup>Julien Toulouse, Roland Assaraf, and Cyrus J. Umrigar. Chapter Fifteen - Introduction to the Variational and Diffusion Monte Carlo Methods. In Philip E. Hoggan and Telhat Ozdogan, editors, *Electron Correlation in Molecules – ab initio Beyond Gaussian Quantum Chemistry*, volume 73 of *Advances in Quantum Chemistry*, pages 285–314. Academic Press, 2016. doi:https://doi.org/10.1016/bs.aiq.2015.07.003. URL <https://www.sciencedirect.com/science/article/pii/S0065327615000386>. ISSN: 0065-3276.
- <sup>51</sup>Junjiro Kanamori. Electron Correlation and Ferromagnetism of Transition Metals. *Progress of Theoretical Physics*, 30(3):275–289, 09 1963. ISSN 0033-068X. doi:10.1143/PTP.30.275. URL <https://doi.org/10.1143/PTP.30.275>.
- <sup>52</sup>M. Bajdich, L. Mitás, G. Drobný, L. K. Wagner, and K. E. Schmidt. Pfaffian pairing wave functions in electronic-structure quantum monte carlo simulations. *Phys. Rev. Lett.*, 96:130201, Apr 2006. doi:10.1103/PhysRevLett.96.130201. URL <https://link.aps.org/doi/10.1103/PhysRevLett.96.130201>.
- <sup>53</sup>Di Luo and Bryan K. Clark. Backflow transformations via neural networks for quantum many-body wave functions. *Phys. Rev. Lett.*, 122:226401, Jun 2019. doi:10.1103/PhysRevLett.122.226401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.122.226401>.
- <sup>54</sup>Jan Hermann, James Spencer, Kenny Choo, Antonio Mezzacapo, W. M. C. Foulkes, David Pfau, Giuseppe Carleo, and Frank Noé. Ab initio quantum chemistry with neural-network wavefunctions. *Nature Reviews Chemistry*, 7(10):692–709, October 2023. ISSN 2397-3358. doi:10.1038/s41570-023-00516-8. URL <https://doi.org/10.1038/s41570-023-00516-8>.
- <sup>55</sup>Z. Schätzle, P. B. Szabó, M. Mezera, J. Hermann, and F. Noé. Deep-QMC: An open-source software suite for variational optimization of deep-learning molecular wave functions. *The Journal of Chemical Physics*, 159(9):094108, 09 2023. ISSN 0021-9606. doi:10.1063/5.0157512. URL <https://doi.org/10.1063/5.0157512>.
- <sup>56</sup>Haotian Ye, Ruichen Li, Yuntian Gu, Yiping Lu, Di He, and Liwei Wang.  $\tilde{O}(n^2)$  representation of general continuous anti-symmetric function, 2024.
- <sup>57</sup>Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017. doi:10.1126/science.aag2302.
- <sup>58</sup>Gabriel Pescia, Jannes Nys, Jane Kim, Alessandro Lovato, and Giuseppe Carleo. Message-passing neural quantum states for the homogeneous electron gas, 2023.





- <sup>59</sup>Remmy Zen, Long My, Ryan Tan, Frédéric Hébert, Mario Gattobigio, Christian Miniatura, Dario Poletti, and Stéphane Bressan. Transfer learning for scalability of neural-network quantum states. *Phys. Rev. E*, 101: 053301, May 2020. doi:10.1103/PhysRevE.101.053301. URL <https://link.aps.org/doi/10.1103/PhysRevE.101.053301>.
- <sup>60</sup>Kenny Choo, Antonio Mezzacapo, and Giuseppe Carleo. Fermionic neural-network states for ab-initio electronic structure. *Nature Communications*, 11:2368, 2020. doi:10.1038/s41467-020-15724-9.
- <sup>61</sup>Yannic Rath and George H. Booth. Quantum gaussian process state: A kernel-inspired state with quantum support data. *Phys. Rev. Res.*, 4:023126, May 2022. doi:10.1103/PhysRevResearch.4.023126. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.4.023126>.
- <sup>62</sup>Volker L. Deringer, Albert P. Bartók, Noam Bernstein, David M. Wilkins, Michele Ceriotti, and Gábor Csányi. Gaussian process regression for materials and molecules. *Chemical Reviews*, 121(16):10073–10141, 2021. doi:10.1021/acs.chemrev.1c00022. URL <https://doi.org/10.1021/acs.chemrev.1c00022>. PMID: 34398616.
- <sup>63</sup>David J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, New York, NY, USA, 2002. ISBN 0521642981.
- <sup>64</sup>Albert P. Bartók and Gábor Csányi. *International Journal of Quantum Chemistry*, 115(16):1051–1057, 2015. doi: <https://doi.org/10.1002/qua.24927>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qua.24927>.
- <sup>65</sup>Albert P. Bartók, Risi Kondor, and Gábor Csányi. *Phys. Rev. B*, 87: 184115, May 2013. doi:10.1103/PhysRevB.87.184115. URL <https://link.aps.org/doi/10.1103/PhysRevB.87.184115>.
- <sup>66</sup>Matthias Rupp. Machine learning for quantum mechanics in a nutshell. *International Journal of Quantum Chemistry*, 115(16):1058–1073, 2015. doi:10.1002/qua.24954. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qua.24954>.
- <sup>67</sup>Lauri Himanen, Marc O.J. Jäger, Eiaki V. Morooka, Filippo Federici Canova, Yashasvi S. Ranawat, David Z. Gao, Patrick Rinke, and Adam S. Foster. *Computer Physics Communications*, 247:106949, 2020. ISSN 0010-4655. doi:<https://doi.org/10.1016/j.cpc.2019.106949>. URL <https://www.sciencedirect.com/science/article/pii/S0010465519303042>.
- <sup>68</sup>Stuart J. Davie, Nicodemo Di Pasquale, and Paul L. A. Popelier. Kriging atomic properties with a variable number of inputs. *The Journal of Chemical Physics*, 145(10):104104, 2016. doi:10.1063/1.4962197. URL <https://doi.org/10.1063/1.4962197>.
- <sup>69</sup>Sandip De, Albert P. Bartók, Gábor Csányi, and Michele Ceriotti. *Phys. Chem. Chem. Phys.*, 18:13754–13769, 2016. doi:10.1039/C6CP00415F. URL <http://dx.doi.org/10.1039/C6CP00415F>.
- <sup>70</sup>Mario Motta, Claudio Genovese, Fengjie Ma, Zhi-Hao Cui, Randy Sawaya, Garnet Kin-Lic Chan, Natalia Chepiga, Phillip Helms, Carlos Jiménez-Hoyos, Andrew J. Millis, Ushnish Ray, Enrico Ronca, Hao Shi, Sandro Sorella, Edwin M. Stoudenmire, Steven R. White, and Shiwei Zhang. *Phys. Rev. X*, 10:031058, Sep 2020. doi:10.1103/PhysRevX.10.031058. URL <https://link.aps.org/doi/10.1103/PhysRevX.10.031058>.
- <sup>71</sup>Anton V. Sinitskiy, Loren Greenman, and David A. Mazziotti. Strong correlation in hydrogen chains and lattices using the variational two-electron reduced density matrix method. *The Journal of Chemical Physics*, 133(1): 014104, 2010.
- <sup>72</sup>Lorenzo Stella, Claudio Attaccalite, Sandro Sorella, and Angel Rubio. Strong electronic correlation in the hydrogen chain: A variational monte carlo study. *Phys. Rev. B*, 84:245117, Dec 2011. doi:10.1103/PhysRevB.84.245117. URL <https://link.aps.org/doi/10.1103/PhysRevB.84.245117>.
- <sup>73</sup>Johannes Hachmann, Wim Cardoen, and Garnet Kin-Lic Chan. Multireference correlation in long molecules with the quadratic scaling density matrix renormalization group. *The Journal of chemical physics*, 125(14):144101, 2006.
- <sup>74</sup>A. Szabo and N.S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Books on Chemistry. Dover Publications, 2012. ISBN 9780486134598. URL <https://books.google.com/books?id=RQ3DgAAQBAJ>.
- <sup>75</sup>Hao Shi and Shiwei Zhang. Some recent developments in auxiliary-field quantum monte carlo for real materials. *The Journal of Chemical Physics*, 154(2):024107, 2021. doi:10.1063/5.0031024. URL <https://doi.org/10.1063/5.0031024>.
- <sup>76</sup>Isaiah Shavitt and Rodney J. Bartlett. *Many-Body Methods in Chemistry and Physics: MBPT and Coupled-Cluster Theory*. Cambridge Molecular Science, Cambridge University Press, 2009. doi:10.1017/CBO9780511596834.
- <sup>77</sup>James McClain, Qiming Sun, Garnet Kin-Lic Chan, and Timothy C. Berkelbach. Gaussian-based coupled-cluster theory for the ground-state and band structure of solids. *Journal of Chemical Theory and Computation*, 13(3):1209–1218, 2017.
- <sup>78</sup>Xiao Wang and Timothy C. Berkelbach. Excitons in solids from periodic equation-of-motion coupled-cluster theory. *Journal of Chemical Theory and Computation*, 16(5):3095–3103, 2020.
- <sup>79</sup>Shiwei Zhang and Henry Krakauer. Quantum monte carlo method using phase-free random walks with slater determinants. *Phys. Rev. Lett.*, 90: 136401, Apr 2003. doi:10.1103/PhysRevLett.90.136401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.90.136401>.
- <sup>80</sup>Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. An auxiliary-field quantum monte carlo study of the chromium dimer. *The Journal of Chemical Physics*, 142(6):064302, 2015.
- <sup>81</sup>W. A. Al-Saidi, Shiwei Zhang, and Henry Krakauer. Bond breaking with auxiliary-field quantum monte carlo. *The Journal of Chemical Physics*, 127 (14):144101, 2007.
- <sup>82</sup>James Shee, Benjamin Rudshyeyn, Evan J. Arthur, Shiwei Zhang, David R. Reichman, and Richard A. Friesner. On achieving high accuracy in quantum chemical calculations of 3d transition metal-containing systems: A comparison of auxiliary-field quantum monte carlo with coupled cluster, density functional theory, and experiment for diatomic molecules. *Journal of Chemical Theory and Computation*, 15(4):2346–2358, 2019.
- <sup>83</sup>Joonho Lee, Fionn D. Malone, and Miguel A. Morales. Utilizing essential symmetry breaking in auxiliary-field quantum monte carlo: Application to the spin gaps of the c36 fullerene and an iron porphyrin model complex. *Journal of Chemical Theory and Computation*, 16(5):3019–3027, 2020.
- <sup>84</sup>Fengjie Ma, Wirawan Purwanto, Shiwei Zhang, and Henry Krakauer. Quantum monte carlo calculations in solids with downfolded hamiltonians. *Phys. Rev. Lett.*, 114:226401, Jun 2015.
- <sup>85</sup>Fionn D. Malone, Anouar Benali, Miguel A. Morales, Michel Caffarel, Paul R. C. Kent, and Luke Shulenburger. Systematic comparison and cross-validation of fixed-node diffusion monte carlo and phaseless auxiliary-field quantum monte carlo in solids. *Phys. Rev. B*, 102:161104, Oct 2020. doi:10.1103/PhysRevB.102.161104. URL <https://link.aps.org/doi/10.1103/PhysRevB.102.161104>.
- <sup>86</sup>Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D. McClain, Elvira R. Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. *WIREs Computational Molecular Science*, 8(1):e1340, 2018. doi:<https://doi.org/10.1002/wcms.1340>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1340>.
- <sup>87</sup>Jeongnim Kim, Andrew D Baczewski, Todd D Beaudet, Anouar Benali, M Chandler Bennett, Mark A Berrill, Nick S Blunt, Edgar Josué Landinez Borda, Michele Casula, David M Ceperley, Simone Chiesa, Bryan K Clark, Raymond C Clay, Kris T Delaney, Mark Dewing, Kenneth P Esler, Hongxia Hao, Olle Heinonen, Paul R C Kent, Jaron T Krogel, Ilkka Kylänpää, Ying Wai Li, M Graham Lopez, Ye Luo, Fionn D Malone, Richard M Martin, Amrita Mathuriya, Jeremy McMinis, Cody A Melton, Lubos Mitas, Miguel A Morales, Eric Neuscamman, William D Parker, Sergio D Pineda Flores, Nichols A Romero, Brenda M Rubenstein, Jacqueline A R Shea, Hyeondeok Shin, Luke Shulenburger, Andreas F Tillack, Joshua P Townsend, Norm M Tubman, Brett Van Der Goetz, Jordan E Vincent, D ChangMo Yang, Yubo Yang, Shuai Zhang, and Luning Zhao. QMCPACK: an open source ab initio quantum monte carlo package for the electronic structure of atoms, molecules and solids. *Journal of Physics: Condensed Matter*, 30(19):195901, apr 2018. doi:10.1088/1361-648x/aab9c3. URL <https://doi.org/10.1088/1361-648x/aab9c3>.
- <sup>88</sup>Giulio Imbalzano, Andrea Anelli, Daniele Giofré, Sinja Klees, Jörg Behler, and Michele Ceriotti. *The Journal of Chemical Physics*, 148(24):241730, 2018. doi:10.1063/1.5024611. URL <https://doi.org/10.1063/1.5024611>.
- <sup>89</sup>Matan Gavish and David L. Donoho. The optimal hard threshold for singular values is  $\frac{4}{\sqrt{3}}$ . *IEEE Transactions on Information Theory*, 60(8):



5040–5053, 2014. doi:10.1109/TIT.2014.2323359.

<sup>90</sup>Rose K Cersonsky, Benjamin A Helfrecht, Edgar A Engel, Sergei Klavinek, and Michele Ceriotti. Improving sample and feature selection with principal covariates regression. *Machine Learning: Science and Technology*, 2(3):035038, jul 2021. doi:10.1088/2632-2153/abfe7c. URL <https://doi.org/10.1088/2632-2153/abfe7c>.

<sup>91</sup>W. A. Al-Saidi, Shiwei Zhang, and Henry Krakauer. Bond breaking with auxiliary-field quantum monte carlo. *The Journal of Chemical Physics*, 127(14):144101, 2007. doi:10.1063/1.2770707. URL <https://doi.org/10.1063/1.2770707>.

<sup>92</sup>Jie Cui and Roman V Krems. Efficient non-parametric fitting of potential energy surfaces for polyatomic molecules with gaussian processes. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(22):224001, oct

2016. doi:10.1088/0953-4075/49/22/224001. URL <https://doi.org/10.1088/0953-4075/49/22/224001>.

<sup>93</sup>Chiraag Nataraj, Edgar Josué Landinez Borda, Axel van de Walle, and Amit Samanta. A systematic analysis of phase stability in refractory high entropy alloys utilizing linear and non-linear cluster expansion models. *Acta Materialia*, 220:117269, 2021. ISSN 1359-6454. doi:<https://doi.org/10.1016/j.actamat.2021.117269>. URL <https://www.sciencedirect.com/science/article/pii/S1359645421006492>.

<sup>94</sup>Hao Shi and Shiwei Zhang. Infinite variance in fermion quantum monte carlo calculations. *Phys. Rev. E*, 93:033303, Mar 2016. doi:10.1103/PhysRevE.93.033303. URL <https://link.aps.org/doi/10.1103/PhysRevE.93.033303>.

