

## Active learning approach to simulations of strongly correlated matter with the ghost Gutzwiller approximation

Marius S. Frank <sup>1</sup>, Denis G. Artiukhin <sup>2</sup>, Tsung-Han Lee,<sup>3</sup> Yongxin Yao <sup>4,5</sup>, Kipton Barros,<sup>6</sup> Ove Christiansen,<sup>7</sup> and Nicola Lanatà <sup>8,9,\*</sup>

<sup>1</sup>Department of Chemistry, Aarhus University, Langelandsgade 140, 8000 Aarhus C

<sup>2</sup>Institut für Chemie und Biochemie, Freie Universität Berlin, Arnimallee 22, 14195 Berlin, Germany

<sup>3</sup>Department of Physics, National Chung Cheng University, Chiayi 62102, Taiwan

<sup>4</sup>Ames National Laboratory, U.S. Department of Energy, Ames, Iowa 50011, USA

<sup>5</sup>Department of Physics and Astronomy, Iowa State University, Ames, Iowa 50011, USA

<sup>6</sup>Theoretical Division and CNLS, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

<sup>7</sup>Department of Chemistry, Aarhus University, Langelandsgade 140, 8000 Aarhus C

<sup>8</sup>School of Physics and Astronomy, Rochester Institute of Technology, 84 Lomb Memorial Drive, Rochester, New York 14623, USA

<sup>9</sup>Center for Computational Quantum Physics, Flatiron Institute, New York, New York 10010, USA



(Received 8 December 2023; revised 8 February 2024; accepted 16 February 2024; published 5 March 2024)

Quantum embedding (QE) methods such as the ghost Gutzwiller approximation (gGA) offer a powerful approach to simulating strongly correlated systems, but come with the computational bottleneck of computing the ground state of an auxiliary embedding Hamiltonian (EH) iteratively. In this work, we introduce an active learning (AL) framework integrated within the gGA to address this challenge. The methodology is applied to the single-band Hubbard model and results in a significant reduction in the number of instances where the EH must be solved. Through a principal component analysis (PCA), we find that the EH parameters form a low-dimensional structure that is largely independent of the geometric specifics of the systems, especially in the strongly correlated regime. Our AL strategy enables us to discover this low-dimensionality structure on the fly, while leveraging it for reducing the computational cost of gGA, laying the groundwork for more efficient simulations of complex strongly correlated materials.

DOI: [10.1103/PhysRevResearch.6.013242](https://doi.org/10.1103/PhysRevResearch.6.013242)

### I. INTRODUCTION

At present, most quantitative simulations of quantum matter utilize standard approximations to density functional theory (DFT) [1,2]. However, these approximations face limitations when simulating the properties of strongly correlated systems, which are solids and molecules where electrons are localized around specific atomic sites and have intensified interactions due to spatial confinement. This issue is especially relevant in materials containing transition metals from the three-dimensional (3D) series and to lanthanides and actinides. To address this challenge, various QE [3,4] many-body techniques have been developed. Methods such as dynamical mean-field theory (DMFT) [5–9], density matrix embedding theory (DMET) [10,11], rotationally invariant slave boson theory (RISB) [12–14], and the multiorbital Gutzwiller approximation (GA) [15–22], are now widely used for quantitatively simulating strongly correlated systems.

Recently an extension of the GA, denoted as the gGA [23,24], has been developed. The gGA framework incorporates auxiliary Fermionic degrees of freedom to enrich the variational space. Notably, gGA has demonstrated accuracy that is comparable to DMFT [23–27], indicating that it might serve as an advantageous alternative, especially when aiming for a combination of accuracy and computational manageability.

However, all of the available QE many-body techniques pose a computational burden for emerging applications in materials discovery, where computational efficiency is crucial for reducing both the time and cost of material development. The main reason lies in their common QE algorithmic structure, that requires the iterative solution of an EH for each correlated fragment in the system, constituting the most computationally intensive step [20,28–30]. Addressing this bottleneck could enable accurate simulations of strongly correlated materials at computational costs comparable to traditional approximations to DFT.

In prior work, a machine-learning-based solution to this problem was proposed both in the context of DMFT [31,32] and in the context of the GA [33], exploiting the observation that the mathematical structure of the EH is determined solely by the electron shell structure of the impurity, thus being consistent across diverse materials and molecules. This intrinsic commonality that arises when solving the EH across different materials and molecules, which we refer to as “universality,”

\*Corresponding author: [nxslps@rit.edu](mailto:nxslps@rit.edu)

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

suggests that machine learning (ML) techniques could, in principle, be trained *once and for all* to solve the EH problem, thereby bypassing the computational bottleneck of all subsequent QE simulations. In particular, in Ref. [33] a method combining the GA method with a ML algorithm, termed “n-KRR,” demonstrated success in implementing this program for a series of actinide systems. However, this achievement was enabled by the possibility to specifically conjecture the physically relevant range of training data for these materials—an advantage that is not generally available. Indeed, this represents the primary barrier to overcome for extending this approach to general many-body systems: it is generally impossible to preemptively determine which training data should be generated. Consequently, a different approach is required to make this ML strategy universally applicable.

To overcome the challenge of determining *a priori* training data, here we introduce an active learning methodology that marries probabilistic ML techniques—specifically, a recent extension of Gaussian process regression (GPR)—with the gGA framework. As new gGA calculations proceed, our active learning model continuously evaluates new instances of the EH problem, adaptively updating and refining its own training set based on the level of uncertainty in its predictions. This strategy eliminates the need for a predefined training set and ensures that only physically relevant data are gathered throughout the computational process.

We benchmark our method using the single-band Hubbard model across varying geometries and interaction strengths, thereby significantly reducing the required number of explicit EH calculations. Using a PCA, we show that the EH parameters explored throughout these calculations have a low-dimensional structure, largely independent of the specific lattice configurations, particularly in strongly correlated regimes. We discuss how such inherent low-dimensional structure of the parameter space opens a path for computational techniques commonly found in computer science, underlining the potential of our active learning strategy to generalize across a wide array of strongly correlated materials in future work.

## II. MODEL AND gGA METHOD

This section aims to lay the foundation for the subsequent development of the QE algorithmic structure. We begin by introducing the single-band Hubbard model, that we employ in our benchmark calculations. We then present the formulation of the gGA. A primary focus of gGA is to iteratively solve for the ground state of an EH, an essential component of the QE approach.

### A. Single-band Hubbard model

For clarity, in this work we present the formalism underlying our AL framework focusing on a generic single-band Hubbard Hamiltonians represented as follows:

$$\hat{H} = \sum_{\substack{i,j=1 \\ i \neq j}}^{\mathcal{N}} \sum_{\sigma=\uparrow,\downarrow} t_{ij} c_{i\sigma}^\dagger c_{j\sigma} + \sum_{i=1}^{\mathcal{N}} \frac{U}{2} (\hat{n}_i - 1)^2 - \mu \sum_{i=1}^{\mathcal{N}} \hat{n}_i, \quad (1)$$

where  $\mathcal{N}$  is the number of system sites (which we refer to as “fragments” from now on),  $c_{i\sigma}^\dagger$  and  $c_{i\sigma}$  are Fermionic creation and annihilation operators,  $\sigma$  is a spin label,  $i$  and  $j$  are fragment labels,  $\mu$  is the chemical potential,  $U$  is the interaction strength and  $\hat{n}_i = \sum_{\sigma} c_{i\sigma}^\dagger c_{i\sigma}$  is the number operator for the system fragment  $i$ , and  $t_{ij}$  is a generic hopping matrix, with arbitrary entries.

### B. gGA Lagrange function

Specializing the theory presented in Refs. [23,24] to the single-orbital Hubbard equation given by Eq. (1), and focusing on solutions preserving both spin and translational symmetries, we find that the ground state in the gGA can be obtained by extremizing the following Lagrange function:

$$\begin{aligned} \mathcal{L}[\Phi, E^c; \mathcal{R}, \Lambda; \mathcal{D}, \Lambda^c; \Delta, \Psi_0, E; \mu] &= \langle \Psi_0 | \hat{H}_{\text{qp}}[\mathcal{R}, \Lambda] | \Psi_0 \rangle + E(1 - \langle \Psi_0 | \Psi_0 \rangle) \\ &+ \sum_{i=1}^{\mathcal{N}} [\langle \Phi_i | \hat{H}_{\text{emb}}^i[\mathcal{D}_i, \Lambda_i^c, U, \mu] | \Phi_i \rangle + E_i^c(1 - \langle \Phi_i | \Phi_i \rangle)] \\ &- \sum_{i=1}^{\mathcal{N}} \left[ \sum_{\sigma=\uparrow,\downarrow} \sum_{a,b=1}^B ([\Lambda_i]_{ab} + [\Lambda_i^c]_{ab}) [\Delta_i]_{ab} \right. \\ &\left. + \sum_{\sigma=\uparrow,\downarrow} \sum_{c,a=1}^B ([\mathcal{D}_i]_a [\mathcal{R}_i]_c [\Delta_i (\mathbb{I} - \Delta_i)]_{ca}^{\frac{1}{2}} + \text{c.c.}) \right], \quad (2) \end{aligned}$$

where  $\mathbb{I}$  is the identity matrix, the integer number  $B$  controls the size of the gGA variational space and, in turn, the precision of the gGA approach,  $E$  and  $E_i^c$  are scalars, and  $\Delta_i$ ,  $\Lambda_i$ , and  $\Lambda_i^c$  are  $B \times B$  Hermitian matrices. Additionally,  $\mathcal{D}_i$  and  $\mathcal{R}_i$  are  $B \times 1$  column matrices. The so-called “quasiparticle Hamiltonian” ( $\hat{H}_{\text{qp}}$ ) and EH ( $\hat{H}_{\text{emb}}^i$ ) are defined as

$$\begin{aligned} \hat{H}_{\text{qp}}[\mathcal{R}, \Lambda] &= \sum_{i=1}^{\mathcal{N}} \sum_{a,b=1}^B \sum_{\sigma=\uparrow,\downarrow} [\Lambda_i]_{ab} f_{i\sigma}^\dagger f_{ib\sigma} \\ &+ \sum_{\substack{i,j=1 \\ i \neq j}}^{\mathcal{N}} \sum_{a,b=1}^B \sum_{\sigma=\uparrow,\downarrow} [\mathcal{R}_i t_{ij} \mathcal{R}_j^\dagger]_{ab} f_{i\sigma}^\dagger f_{jb\sigma}, \quad (3) \\ \hat{H}_{\text{emb}}^i[\mathcal{D}_i, \Lambda_i^c, U, \mu] &= \frac{U}{2} (\hat{n}_i - 1)^2 - \mu \hat{n}_i \\ &+ \sum_{a=1}^B \sum_{\sigma=\uparrow,\downarrow} [[\mathcal{D}_i]_a c_{i\sigma}^\dagger b_{i\sigma} + \text{H.c.}] \\ &+ \sum_{a,b=1}^B \sum_{\sigma=\uparrow,\downarrow} [\Lambda_i^c]_{ab} b_{ib\sigma} b_{i\sigma}^\dagger. \quad (4) \end{aligned}$$

Here the vector  $|\Phi_i\rangle$  is the most general embedding state for the fragment  $i$ , i.e., the most general state within the Fock space spanned by the  $2(B+1)$  modes  $c_{i\sigma}^\dagger$  and  $b_{i\sigma}^\dagger$ , with  $(B+1)$  Fermions in total (half-filled). The vector  $|\Psi_0\rangle$  is the most general single-particle state within the so-called “quasiparticle” space, spanned by the  $2B\mathcal{N}$  modes  $f_{i\sigma}^\dagger$ .

For  $B=1$ , Eq. (2) reduces to the standard GA Lagrange function. In this work we set  $B=3$ , which proved to be sufficient for capturing the ground-state properties with

accuracy comparable to DMFT for the ground-state properties [23–27].

### C. Gauge invariance and physical observables

It can be readily verified that the gGA Lagrangian is invariant with respect to the following gauge transformation:

$$|\Psi_0\rangle \rightarrow \mathcal{U}^\dagger(\theta)|\Psi_0\rangle, \quad (5)$$

$$|\Phi_i\rangle \rightarrow U_i^\dagger(\theta)|\Phi_i\rangle, \quad (6)$$

$$\mathcal{R}_i \rightarrow u_i^\dagger(\theta)\mathcal{R}_i, \quad (7)$$

$$\mathcal{D}_i \rightarrow u_i^T(\theta)\mathcal{D}_i, \quad (8)$$

$$\Delta_i \rightarrow u_i^T(\theta)\Delta_i u_i^*(\theta), \quad (9)$$

$$\Lambda_i \rightarrow u_i^\dagger(\theta)\Lambda_i u_i(\theta), \quad (10)$$

$$\Lambda_i^c \rightarrow u_i^\dagger(\theta)\Lambda_i^c u_i(\theta), \quad (11)$$

with

$$u_i(\theta_i) = e^{i\theta_i}, \quad (12)$$

$$U_i(\theta_i) = e^{i\sum_{a,b=1}^B[\theta_i]_{ab}b_{ia}^\dagger b_{ib}}, \quad (13)$$

$$\mathcal{U}(\theta) = e^{i\sum_{i=1}^{\mathcal{N}}\sum_{a,b=1}^B[\theta_i]_{ab}f_{ia}^\dagger f_{ib}}, \quad (14)$$

where  $\theta = (\theta_1, \dots, \theta_{\mathcal{N}})$ ,  $\theta_i$  are  $B \times B$  Hermitian matrices and the superscript "T" denotes the transpose, while the superscript "\*" denotes the complex conjugate. The name "gauge" here refers to the fact that modifications of the parameters generated by such a gauge transformation do not influence any physical observable, which can be extracted from the variational parameters that extremize the Lagrange function.

For completeness, below we write explicitly how the physical observables that we calculate in this paper are computed as a function of the variational parameters, based on the theoretical framework derived in previous work [23,24].

The total energy of the system is given by the Lagrange function value after extremization (which is gauge invariant). The expectation values for local observables are encoded in  $|\Phi_i\rangle$ . In particular, the local double-occupancy expectation value in the gGA ground state is given by the following gauge-invariant expression:

$$\langle \hat{n}_{i\uparrow}\hat{n}_{i\downarrow} \rangle_{\text{gGA}} = \langle \Phi_i | \hat{n}_{i\uparrow}\hat{n}_{i\downarrow} | \Phi_i \rangle,$$

where  $\hat{n}_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ . To calculate the quasiparticle weight, it is convenient to express the variational parameters in a gauge where  $[\Lambda_i]_{ab} = [l_i]_a \delta_{ab}$  (which always exists, since  $\Lambda_i$  is Hermitian). In this gauge, the mathematical expression for the quasiparticle weight is the following:

$$Z_i = \left[ 1 - \frac{\partial \Sigma_i}{\partial \omega} \right]_{\omega=0}^{-1} = \frac{([l_i]_2 [l_i]_3 [\mathcal{R}_i]_1^2 + [l_i]_1 [l_i]_3 [\mathcal{R}_i]_2^2 + [l_i]_1 [l_i]_2 [\mathcal{R}_i]_3^2)}{[l_i]_1^2 [l_i]_3^2 [\mathcal{R}_i]_2^2 + [l_i]_2^2 ([l_i]_3^2 [\mathcal{R}_i]_1^2 + l_i^2 [\mathcal{R}_i]_3^2)}, \quad (15)$$

where  $\Sigma_i(\omega)$  is the self-energy.

## III. FORMULATION OF THE ML PROBLEM

A pivotal insight at the base of the ML approach proposed in this work is that the problem of extremizing the Lagrange function in Eq. (2) can be formally tackled by first solving for  $|\Phi_i\rangle$  and  $E_i^c$  as a function of the other parameters. This amounts to replacing the original variational problem of extremizing  $\mathcal{L}$  in Eq. (2) with the problem of extremizing the following Lagrange function:

$$\begin{aligned} \tilde{\mathcal{L}}[\mathcal{R}, \Lambda; \mathcal{D}, \Lambda^c; \Delta, \Psi_0, E; \mu] &= \langle \Psi_0 | \hat{H}_{\text{qp}}[\mathcal{R}, \Lambda] | \Psi_0 \rangle + E(1 - \langle \Psi_0 | \Psi_0 \rangle) \\ &+ \sum_{i=1}^{\mathcal{N}} \bar{E}^c(\mathcal{D}_i, \Lambda_i^c, U, \mu) \\ &- \sum_{i=1}^{\mathcal{N}} \left[ \sum_{\sigma=\uparrow,\downarrow} \sum_{a,b=1}^B ([\Lambda_i]_{ab} + [\Lambda_i^c]_{ab}) [\Delta_i]_{ab} \right. \\ &\left. + \sum_{\sigma=\uparrow,\downarrow} \sum_{c,a=1}^B ([\mathcal{D}_i]_{a[\mathcal{R}_i]_c} [\Delta_i (\mathbb{I} - \Delta_i)]_{ca}^{\frac{1}{2}} + \text{c.c.}) \right], \end{aligned} \quad (16)$$

where

$$\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu) = \langle \hat{H}_{\text{emb}}^i[\mathcal{D}, \Lambda^c, U, \mu] \rangle_{\mathcal{D}, \Lambda^c, U, \mu}, \quad (17)$$

$\hat{H}_{\text{emb}}^i$  is the EH defined in Eq. (4), and the expectation value is taken with respect to the corresponding half-filled ground state  $|\Phi_i\rangle$ . Note that the subscript  $i$  is not present in the function  $\bar{E}^c$  of Eq. (16), highlighting the fact that the function does *not* depend on it.

A key property of the function  $\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu)$  is that it is "universal," in the sense that its definition is irrespective of the details of the model system under study, but it only depends on the type of fragment (e.g., whether it corresponds to a shell of s,p,d or f electrons) and the gGA variational space employed (i.e., the parameter  $B$  used). For example, when applying gGA to any of the Hubbard models defined in Eq. (45), the form of  $\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu)$  would remain consistent, independent of the specific lattice structure or the numerical values of the hopping matrix  $t$ . Therefore, if one could learn the energy function  $\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu)$ , along with its gradient, then the computational cost of extremizing  $\tilde{\mathcal{L}}$  would be significantly alleviated. The purpose of this work is to derive a ML model to accomplish this goal.

### A. EH Universal Function

The Lagrange equations obtained by extremizing the Lagrange function in Eq. (16) can be solved iteratively [23,24]. The specific algorithm employed in this work is detailed in the Supplemental Material [34].

For the purposes of this paper, the essential point is that the algorithmic structure involves iteratively evaluating  $\bar{E}^c(\mathcal{D}_i, \Lambda_i^c, U, \mu)$  and its gradient, as the parameters  $\mathcal{D}_i$  and  $\Lambda_i^c$  are updated at each step. Specifically, the computational

bottleneck lies in the evaluation of the ground-state single-particle density matrix elements of each fragment  $i$ :

$$\frac{\partial \bar{E}^c}{\partial [\mathcal{D}_i]_{ia}} = 2[\rho_i^{\text{hyb}}]_a = 2 \left\langle \sum_{\sigma} c_{i\sigma}^{\dagger} b_{i\sigma} \right\rangle_{\mathcal{D}_i, \Lambda_i^c, U, \mu}, \quad (18)$$

$$\frac{\partial \bar{E}^c}{\partial [\Lambda_i^c]_{ab}} = [\rho_i^{\text{bath}}]_{ab} = \left\langle \sum_{\sigma} b_{i\sigma} b_{i\sigma}^{\dagger} \right\rangle_{\mathcal{D}_i, \Lambda_i^c, U, \mu}, \quad (19)$$

where the identities above hold true because of the Hellmann–Feynman theorem.

### B. Reducing the complexity of the learning problem

In the previous subsection we have introduced the function  $\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu)$ . Note that, since  $\mathcal{D}$  and  $\Lambda^c$  are generally complex and  $\Lambda^c$  is Hermitian,  $\bar{E}^c$  is a function of  $2B + B^2 + 2$  real parameters.

In this section we show that it is possible to reduce the problem of learning the aforementioned universal EH energy function to the problem of learning the following function of only  $2B$  real variables:

$$\begin{aligned} \mathcal{E}(\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c) \\ = \langle \hat{H}_{\text{emb}}[\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c] \rangle_{\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c}, \end{aligned} \quad (20)$$

representing the ground-state energy of the following Hamiltonian:

$$\begin{aligned} \hat{H}_{\text{emb}}[\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c] \\ = \frac{1}{2}(\hat{n} - 1)^2 + \sum_{a=1}^B \sum_{\sigma=\uparrow, \downarrow} [ \tilde{\mathcal{D}}_a c_{\sigma}^{\dagger} b_{a\sigma} + \text{H.c.} ] \\ + \sum_{a=1}^B \sum_{\sigma=\uparrow, \downarrow} \tilde{\Lambda}_{aa}^c b_{a\sigma} b_{a\sigma}^{\dagger} \end{aligned} \quad (21)$$

at half filling. Furthermore, it is sufficient to learn the function  $\mathcal{E}$ , defined by Eq. (20), over the restricted domain such that  $\tilde{\mathcal{D}}_1 \geq \tilde{\mathcal{D}}_2 \geq \dots \geq \tilde{\mathcal{D}}_B$ .

This simplification is made possible by the fact that the EH function  $\bar{E}^c$ , defined in Sec. III A, satisfies the following general properties:

(1) Invariance of half-filled ground state  $|\Phi\rangle$  under simultaneous shift of impurity and bath energies:

$$\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu) = \bar{E}^c(\mathcal{D}, \Lambda^c + \mu \mathbb{I}, U, 0) - \mu, \quad (22)$$

(2) Linear homogeneity:

$$\bar{E}^c(x\mathcal{D}, x\Lambda^c, xU, x\mu) = x\bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu) \quad \forall x, \quad (23)$$

(3) Gauge invariance:

$$\bar{E}^c(u^T \mathcal{D}, u^{\dagger} \Lambda^c u, U, \mu) = \bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu) \quad \forall u, \quad (24)$$

where  $x$  is any real number and  $u$  is any  $B \times B$  unitary matrix (i.e., a gauge transformation, see Sec. II C). From the properties above it follows that

$$\begin{aligned} \bar{E}^c(\mathcal{D}, \Lambda^c, U, \mu) &= \bar{E}^c(\mathcal{D}, \Lambda^c + \mu \mathbb{I}, U, 0) + \mu \\ &= U \bar{E}^c(\tilde{\mathcal{D}}, \tilde{\Lambda}^c, 1, 0) + \mu \\ &= U \mathcal{E}(\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c) + \mu, \end{aligned} \quad (25)$$

where

$$\tilde{\Lambda}^c = \frac{1}{U} u^{\dagger} (\Lambda^c + \mu \mathbb{I}) u, \quad (26)$$

$$\tilde{\mathcal{D}} = \frac{1}{U} u^T \mathcal{D}, \quad (27)$$

and the unitary matrix  $u$  is a gauge transformation chosen in such a way that  $\tilde{\Lambda}^c$  is diagonal and the entries of  $\tilde{\mathcal{D}}$  are real and sorted in descending order, as detailed in the Supplemental Material [34].

Another important consequence of Eqs. (23) and (24) is that, as shown in the Supplemental Material [34], also the gradient of  $\bar{E}^c$  is fully encoded in the gradient of  $\mathcal{E}$ , which is given by the following equations:

$$\frac{\partial \mathcal{E}}{\partial \tilde{\mathcal{D}}_a} = 2\tilde{\rho}_a^{\text{hyb}} = 2 \left\langle \sum_{\sigma} c_{\sigma}^{\dagger} b_{a\sigma} \right\rangle_{\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c}, \quad (28)$$

$$\frac{\partial \mathcal{E}}{\partial \tilde{\Lambda}_{aa}^c} = \tilde{\rho}_{aa}^{\text{bath}} = \left\langle \sum_{\sigma} b_{a\sigma} b_{a\sigma}^{\dagger} \right\rangle_{\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c}, \quad (29)$$

where the expectation values are taken with respect to the ground state of the EH defined in Eq. (21).

### C. Summary

In summary, in this section we have reduced the problem of solving the gGA equations to an iterative procedure which consists of evaluating iteratively the functions

$$\mathcal{E}(\mathbf{X}) = \mathcal{E}(\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c), \quad (30)$$

$$\mathbf{F}(\mathbf{X}) = \nabla \mathcal{E}(\mathbf{X}), \quad (31)$$

where we have introduced the  $2B$ -dimensional real vector,

$$\mathbf{X} = (\tilde{\mathcal{D}}_1, \dots, \tilde{\mathcal{D}}_B, \tilde{\Lambda}_{11}^c, \dots, \tilde{\Lambda}_{BB}^c). \quad (32)$$

Each evaluation of  $\mathcal{E}(\mathbf{X})$  and  $\mathbf{F}(\mathbf{X})$  requires to calculate the ground state of the EH defined in Eq. (21), whose dimension is  $2^{2(B+1)}$ . While the cost of this operation is essentially negligible for  $B = 3$ , it can quickly become the computational bottleneck of the gGA framework as  $B$  or the fragment's dimension increase, which we aim to mitigate using ML.

An important point to highlight is the simplification introduced by reducing the problem to the learning of a single scalar function  $\mathcal{E}(\mathbf{X})$ , which we achieved through the use of the Hellmann–Feynman theorem. This approach reduces computational cost and improves predictive accuracy compared to the method of individually learning each entry of the ground-state single-particle density matrix of the EH (which is the technique we previously employed in Ref. [33]). First, learning a single scalar function is computationally less demanding than learning an array of functions corresponding to each matrix element. Second, this method also automatically enforces inherent prior information about these specific functions, such as the condition that

$$\nabla \times \mathbf{F}(\mathbf{X}) = 0, \quad (33)$$

i.e.,  $\mathbf{F}(\mathbf{X})$  is conservative, thereby enhancing the overall predictive accuracy of the model.

Interestingly, from the mathematical perspective, the problem outlined above bears a striking resemblance to a specific

successful ML application, namely learning force fields for accelerating molecular dynamics simulations [35–37]. In fact, in both instances, the challenge revolves around learning the total energy  $\mathcal{E}(\mathbf{X})$  of a system and its gradient  $\mathbf{F}(\mathbf{X})$  from computational data. However, there are two key differences:

(1) *Universality*: Within our framework, if we could learn once and for all  $\mathcal{E}(\mathbf{X})$ , then the resulting model could be used to bypass such computational bottleneck for *any* system involving single-orbital fragments (e.g., all models of the form represented in Eq. (45), for all hopping matrices  $t_{ij}$  and for any values of  $U$  and  $\mu$ ).

(2) *Domain structure*: In gGA the sequence of points  $\mathbf{X}_\alpha$  explored throughout the computation of any given model always converges towards the specific  $\bar{\mathbf{X}}$  realizing the corresponding solution. Therefore, the majority of these points, where  $\mathcal{E}(\mathbf{x})$  needs to be learned, gravitate around  $\bar{\mathbf{X}}$ , instead of being spread around the whole parameters space.

These unique characteristics suggest a more tractable learning problem. Specifically, we are not compelled to learn the universal function  $\mathcal{E}(\mathbf{X})$  over its entire domain, hereafter referred to as the “ambient space.” Rather, we can confine our attention to a subset of parameters situated in the proximity of the ground states of physical models. These ground states encapsulate the possible physical embeddings, i.e., the feasible interactions that fragments can have with their environment in the ground state of physical systems. Such parameters presumably constitute a limited fraction of the ambient space.

Therefore, we are confronted with a dual challenge. The first is learning the function  $\mathcal{E}(\mathbf{X})$  over this restricted domain, hereafter referred to as the “latent space.” The second is concurrently unveiling the structure of this latent space, which is expected to have lower dimensionality than the ambient space. The successful completion of the latter task would considerably streamline the overall learning problem, specifically by potentially mitigating the effects of the so-called “curse of dimensionality,” where computational cost grows exponentially with the number of dimensions.

In the next section we describe an AL framework for learning  $\mathcal{E}(\mathbf{X})$  and  $\mathbf{F}(\mathbf{X})$ , which is specifically tailored for this purpose, capitalizing on the observations above.

#### IV. ACTIVE LEARNING FRAMEWORK

In this section, we outline an AL strategy based on probabilistic ML to overcome the computational challenges delineated in the previous section.

A critical aspect of our strategy is the use of a probabilistic ML model, i.e., a model capable of combining our prior knowledge and observed data to make predictions for  $\mathcal{E}$  and its gradient  $\mathbf{F}$  as expectation values  $\langle \mathcal{E}(\mathbf{X}) \rangle$  and  $\langle \mathbf{F}(\mathbf{X}) \rangle$  with respect to a suitable probability distribution, and to quantify their uncertainties through standard deviations:

$$\Sigma_0(\mathbf{X}) = [\langle \mathcal{E}^2(\mathbf{X}) \rangle - \langle \mathcal{E}(\mathbf{X}) \rangle^2]^{\frac{1}{2}}, \quad (34)$$

$$\Sigma_i(\mathbf{X}) = [ \langle F_i^2(\mathbf{X}) \rangle - \langle F_i(\mathbf{X}) \rangle^2 ]^{\frac{1}{2}}, \quad (35)$$

where  $F_i(\mathbf{X})$  are the components of  $\mathbf{F}(\mathbf{X})$ , with  $i = 1, \dots, 2B$ , where  $2B$  is the dimension of  $\mathbf{X}$ . Furthermore, we require a model capable of learning both the energy function  $\mathcal{E}$  and its gradient  $\mathbf{F}$  simultaneously, ensuring exactly the consistency

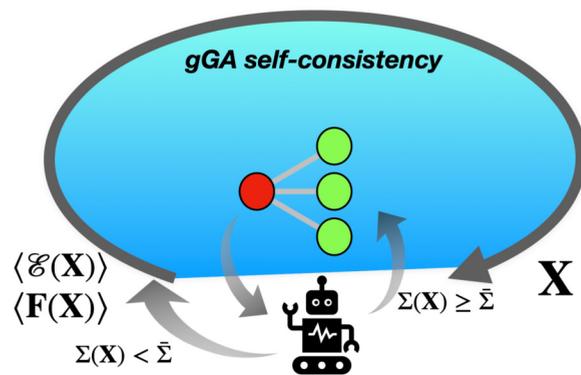


FIG. 1. Schematic representation of the AL structure. A machine is trained on the fly while performing gGA calculations. If the uncertainty estimate  $\Sigma$  for ML predictions at a given point  $\mathbf{X}$  is within a threshold, then the prediction is accepted. Otherwise, the machine database is updated with energy  $\mathcal{E}$  and gradient  $\mathbf{F}$  data, and a new machine evaluation is performed.

between these two quantities. We achieve this by a recent generalization of GPR, satisfying both of these requirements [38,39], described in Sec. IV B and in the Supplemental Material [34].

#### A. gGA + AL Algorithmic Structure

Given a probabilistic ML method with the requirements listed above in place, we proceed to outline the AL strategy as follows:

(1) For each self-consistency cycle, the gGA algorithm produces a parameter vector  $\mathbf{X}$  for every impurity in the system.

(2) The probabilistic ML framework predicts the expectation values  $\langle \mathcal{E}(\mathbf{X}) \rangle$  and  $\langle \mathbf{F}(\mathbf{X}) \rangle$ . It also estimates the uncertainty quantification of these predictions, termed  $\Sigma_0(\mathbf{X})$  and  $\Sigma_i(\mathbf{X})$ , respectively.

(3) If the existing data can produce a sufficiently accurate prediction for  $\langle \mathcal{E}(\mathbf{X}) \rangle$  and  $\langle \mathbf{F}(\mathbf{X}) \rangle$  (with respect to pre-established accuracy thresholds), then the ML model’s outputs are directly employed, thereby bypassing the need for ground-state computation of the EH.

(4) If the data are insufficient for an accurate prediction corresponding to the EH at the given  $\mathbf{X}$ , then new training data are computed for suitable parameters  $\mathbf{X}_\alpha$ . These data points are then added to the database for future use.

(5) The output, obtained through either step 3 or 4, is fed back into the QE algorithm to initiate the next iteration. The process is repeated until convergence is achieved.

The salient features of this framework, as depicted in Fig. 1, are twofold: (1) Since all data are generated through actual gGA calculations, they are inherently situated in the proximity of the “latent space” of physically meaningful embeddings. This circumvents the issue of requiring prior knowledge of the QE latent space, as discussed in Sec. III C. (2) Owing to the universality of the EH, data collected during any calculation for strongly correlated systems are retained permanently. Consequently, the ML component of our framework becomes increasingly efficient with each new calculation, potentially enhancing the computational efficiency of all subsequent gGA simulations.

The details about the algorithm implementation outlined above are described in the subsections below.

### B. Generalized Gaussian process regression

Above in this section we introduced the need for a probabilistic ML model that can learn and predict both the energy function  $\mathcal{E}$  and its gradient  $\mathbf{F}$ , as well as quantify their uncertainties. To meet these requirements, we employ a generalized form of Gaussian process regression (GPR) [38,39], as implemented in a development version of the program package *MidasCpp* [40]. The method constructs a so-called "posterior probability distribution" for the function to learn, which is based on a prior probability distribution and available data.

(1) The prior encodes our expectations about the general properties of the function we aim to learn, such as its range and smoothness.

(2) The observed data in our case consist of a database:

$$\mathbf{D} = \{\mathbf{X}_\alpha, \mathcal{E}_\alpha, \mathbf{F}_\alpha, \sigma_{0\alpha}, \boldsymbol{\sigma}_\alpha\}, \quad (36)$$

where  $\mathbf{X}_\alpha$  are points with evaluated energies  $\mathcal{E}_\alpha$  and gradients  $\mathbf{F}_\alpha$ ,  $\sigma_{0\alpha}$  measures the uncertainty associated with the energy data  $\mathcal{E}_\alpha$ , and  $\boldsymbol{\sigma}_\alpha = (\sigma_{1\alpha}, \dots, \sigma_{d\alpha})$ , where  $\sigma_{i\alpha}$  is the uncertainty associated with the  $i$ -th component of the gradient data  $\mathbf{F}_{i\alpha}$  and  $d$  is the dimension of  $\mathbf{X}_\alpha$ .

In GPR the prior probability distribution is assumed to be a zero-mean Gaussian. Consequently, it is fully characterized by the so-called "kernel function," which is essentially the correlation function  $\langle \mathcal{E}(\mathbf{X})\mathcal{E}(\mathbf{X}') \rangle_{\text{prior}}$ , where the expectation value is taken with respect to the prior probability distribution. The kernel function we employ in this work is the "square exponential kernel," given by

$$\begin{aligned} k(\mathbf{X}, \mathbf{X}') &= \langle \mathcal{E}(\mathbf{X})\mathcal{E}(\mathbf{X}') \rangle_{\text{prior}} \\ &= \sigma_f^2 \exp\left(-\frac{(\mathbf{X} - \mathbf{X}')^2}{2l^2}\right). \end{aligned} \quad (37)$$

This choice is motivated by its simplicity and effectiveness in describing the essential qualitative features of the EH energy landscape that we aim to imbue within the model, which are governed by two hyperparameters:  $l$  and  $\sigma_f$ . The parameter  $l$  is essentially a correlation length, defining the expected "minimum wavelength" or smoothness of the function  $\mathcal{E}(\mathbf{X})$ . The parameter  $\sigma_f$  specifies the expected range of the function, serving as an infrared cutoff. Specifically,  $k(\mathbf{X}, \mathbf{X}) = \sigma_f^2$  corresponds to the expected variance  $\langle \mathcal{E}(\mathbf{X})^2 \rangle_{\text{prior}}$  of the function.

The posterior probability distribution thus integrates both our prior knowledge and the available observed data, generating predictions for  $\mathcal{E}$  and  $\mathbf{F}$  that align with the observed data, while utilizing the prior for making predictions elsewhere. In particular, our AL framework requires to compute the following quantities at any given point  $\mathbf{X}$ :

$$\bar{\mathcal{E}}(\mathbf{X}) = \langle \mathcal{E}(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}}, \quad (38)$$

$$\bar{F}_i(\mathbf{X}) = \langle F_i(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}}, \quad (39)$$

$$\Sigma_0(\mathbf{X}) = \left[ \langle \mathcal{E}^2(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}} - \langle \mathcal{E}(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}}^2 \right]^{\frac{1}{2}}, \quad (40)$$

$$\Sigma_i(\mathbf{X}) = \left[ \langle F_i^2(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}} - \langle F_i(\mathbf{X}) \rangle_{l, \sigma_f, \mathbf{D}}^2 \right]^{\frac{1}{2}}, \quad (41)$$

where the expectation values are taken with respect to the posterior distribution, that depends on the hyperparameters  $l$ ,  $\sigma_f$ , as well as the available data  $\mathbf{D}$ . Explicit expressions for these distributions are provided in the Supplemental Material [34].

It is important to note that, as opposed to standard GPR and the KRR-based method previously used in Ref. [33], the generalized GPR framework outlined above enforces by construction the condition

$$\bar{\mathbf{F}}(\mathbf{X}) = \nabla \bar{\mathcal{E}}(\mathbf{X}), \quad (42)$$

where  $\bar{F}_i(\mathbf{X})$  are the components of  $\bar{\mathbf{F}}(\mathbf{X})$ , with  $i = 1, \dots, d$ , where  $d = 2B$  is the dimension of  $\mathbf{X}$ . Enforcing exactly these conditions yields more accurate predictions. However, as explained in the Supplemental Material [34], such gain comes with additional computational cost. Specifically, if the database  $\mathbf{D}$  contains  $N$  training data points, then making predictions requires to invert a matrix (the so-called "covariance matrix"), whose size is  $N(d+1) \times N(d+1)$ , while it is only  $N \times N$  in standard GPR. In light of this, the computational complexity and matrix size present two challenges that need to be carefully addressed:

(1) **RAM Storage:** The large covariance matrix, of size  $N(d+1) \times N(d+1)$ , necessitates considerable memory storage. This can become a significant issue as the number of data loaded in the GPR framework grows.

(2) **Scalability:** The matrix inversion operation itself has a time complexity of  $\mathcal{O}((N(d+1))^3)$ . As  $N$  increases, it becomes computationally burdensome. Specific measures must be incorporated into our active learning framework to ensure its scalability for large databases.

(3) **Numerical Stability:** The large size of the covariance matrix and closely spaced data points in  $\mathbf{D}$  can make the matrix inversion prone to numerical issues. Specifically, the matrix can become ill-conditioned, having a large condition number, which is the ratio of the largest to the smallest eigenvalue. This makes the matrix sensitive to small changes, potentially leading to a loss of numerical precision.

In the following subsection, we detail how these challenges are addressed in our active learning framework, and also elaborate on our procedure for choosing the hyperparameters  $l$  and  $\sigma_f$ .

### C. Addressing computational and numerical challenges in AL framework

To cope with the computational and numerical challenges presented by the large covariance matrix, we have designed the following strategy within our AL framework.

Our approach is based on the construction of a hyper cubic lattice discretization of the  $d$ -dimensional parameter space  $\mathbf{X}$ , yielding a set of grid points  $\mathbf{X}_{i_1, \dots, i_d}$ . Each  $i_\alpha$  index covers all integers, effectively tiling the whole parameter space. We use a lattice spacing  $a = 0.0125$ , which is set to be much smaller than 1, taking into account the dimensionless nature of our parameters. When predicted values for a given  $\mathbf{X}$  are being requested during the self-consistent embedding computations we find the nearest point on the lattice to this point. In addition we also consider this lattice points  $2d$  nearest neighbors on the lattice (corresponding to  $i_d \pm 1$  for all  $d$ ). Should any of

these  $2d + 1$  points not already be in the database  $\mathbf{D}$ , they are calculated and added. Note that these calculations are independent of each other and, therefore, can be executed in parallel. For each evaluation at  $\mathbf{X}$ , the GPR prediction and its associated uncertainties are calculated based solely on these  $2d + 1$  data points. We locate these nearest points using a k-d tree-based algorithm to maintain computational efficiency. This step effectively imposes a "budget" on the number of training data we use for making GPR predictions, thereby controlling the size of the covariance matrix and the associated computational complexity.

Next, we address the setting of the hyperparameter  $\sigma_f$  within this localized framework. For the GPR prediction of each point  $\mathbf{X}$ , the hyperparameter  $\sigma_f$  for the squared exponential kernel is set based on the data as follows:

$$\sigma_f^2 = \frac{1}{2d + 1} \sum_{\alpha=1}^{2d+1} \mathcal{E}^2(\mathbf{X}_\alpha), \quad (43)$$

which is consistent with its interpretation in terms of the prior probability distribution:  $\sigma_f^2 = \langle \mathcal{E}^2(\mathbf{X}) \rangle_{\text{prior}}$ . To determine the appropriate value of the hyperparameter  $l$ , we employ the following iterative approach:

(1) Initialize  $l$  at  $l_{\text{init}} = 0.5$ . The range for  $l$  is predetermined to be between  $l_{\text{init}}$  and  $l_{\text{final}} = 2.0$ , with increments of  $\Delta l = 0.1$ .

(2) For the current  $l$ , calculate  $\Sigma_{\text{max}}$ , which is the maximum of the uncertainties  $\Sigma_i(\mathbf{X})$  associated with all components of the gradient  $\mathbf{F}(\mathbf{X})$  for  $i = 1, \dots, d$ .

(3) If  $\Sigma_{\text{max}} < \bar{\Sigma} = 10^{-3}$ , then accept the current GPR prediction for that  $l$  and terminate the loop.

(4) Otherwise, increment  $l$  by  $\Delta l$  and return to step 2. If  $l$  exceeds  $l_{\text{final}}$ , then revert to exact calculations for that specific test point and terminate the loop.

The uncertainty parameters  $\sigma_{i\alpha}$ , see Eq. (36), have been all set to  $10^{-5}$  in all of our calculations.

The rationale underlying this algorithm is to initiate with the smallest  $l$  value, thereby making the least assumptions about the landscape's smoothness and relying more heavily on the data for our predictions. If this  $l$  proves insufficient, then we then increment  $l$  in a stepwise manner, each time reassessing the prediction quality. This iterative fine-tuning of  $l$  is in essence a method for optimizing the trade-off between bias and variance, a standard criterion in ML, which ensures that our model is neither too simplistic (high bias; large  $l$ ) nor too sensitive to fluctuations in the data (high variance; small  $l$ ). Note also that our choice of  $l$  range is in line with the fact that our parameters  $\mathbf{X}$  are dimensionless, from which we expect the optimal  $l$  to be around 1. Furthermore, it is consistent with the assumption  $l \gg a$ , acknowledging that we cannot resolve scales smaller than the lattice spacing  $a$ . Finally, note that our choice of  $\bar{\Sigma} = 10^{-3}$  is consistent with the assumption  $\bar{\Sigma} \gg \sigma_{i\alpha}$ , acknowledging that it is impossible to make predictions with higher accuracy than the available data. Thus, the framework above effectively manages challenges related to RAM storage, computational scalability, and numerical stability, while maintaining a balance between computational cost and prediction accuracy.

## V. RESULTS

In this section we document the results obtained by applying the AL algorithm described above to gGA calculations on the single-band Hubbard model at half filling. Here, we consider different values for the Hubbard interaction strength  $U$  and hopping parameters  $t_{ij}$  (corresponding to multiple lattice geometries).

We carry out our calculations within the gGA framework set at  $B = 3$ , which, as proven in previous work, is sufficient for achieving accuracy comparable to DMFT for ground-state properties. As clarified in Sec. III C, in this setting the dimension of the "ambient space" of parameters  $\mathbf{X}$ , where the energy function  $\mathcal{E}(\mathbf{X})$  is defined, is  $2B = 6$ . All training-data evaluations for  $\mathcal{E}(\mathbf{X})$  and  $\mathbf{F}(\mathbf{X})$  are calculated using the exact diagonalization (ED) method.

In Secs. VB and VC we document the efficiency and accuracy obtained in these calculations using our AL approach. In Sec. VD, we demonstrate that the parameters explored span a low-dimensional latent space and discuss the practical implications of these results, as well as their physical interpretation in relation to Mott physics. Additional calculations and analysis for the Hubbard model away from half filling are presented in the Supplemental Material [34].

### A. Goal of benchmark calculations

Our aim is to test the AL method within all interaction regimes of the half-filled Hubbard model at zero temperature, including the so-called coexistence region, which is an interval of parameters  $U$  featuring a metastable Mott state. To capture all of these regimes, all calculations are organized into series constructed as follows. Each series starts from a large value of interaction strength  $U_{\text{max}}$  and decreases it in intervals of  $\Delta U$  to a small value  $U_{\text{min}}$ . Then, the interaction strength is increased back to  $U_{\text{max}}$  with the same spacing. From now on, we refer to such a series of calculations as a "sweep."

A critical metric for quantifying the efficiency of our gGA + AL approach is the ratio of the number of times new data must be acquired and added to the database during a given calculation,  $N_{\text{data}}$ , to the total number of gGA iterations necessary to perform the same calculation without ML,  $N_{\text{iterations}}$ :

$$S = \frac{N_{\text{data}}}{N_{\text{iterations}}}, \quad (44)$$

which we would like to be as small as possible.

It is important to note that the value of  $S$  is heavily influenced by the choice of hyper-parameters, as detailed in Sec. IV C. In these benchmarks we strived for high accuracy by requiring that both the energy  $\mathcal{E}$  and its gradient  $\mathbf{F}$  are estimated to a precision of at least  $\bar{\Sigma} = 10^{-3}$ . Furthermore, a minimum of  $2d + 1$  training points are required within a grid with tight lattice spacing  $a = 0.0125$  for each test point. This ensures that new calculations are invoked when the exploration enters a new parameter region spaced by more than  $a$ .

In the forthcoming benchmark calculations, we aim to address three specific scientific questions for evaluating the utility and efficiency of our gGA + AL method:

(1) **Ability to learn:** Can a sweep of gGA + AL calculations, once completed and with the data stored, be repeated

without requiring any new data for the Hamiltonian parameters already explored? This is a necessary condition for realizing the computational benefits of our data-driven approach.

(2) **Transfer-learning efficiency:** If multiple sweeps are performed, each with different settings such as  $\Delta U$  or  $t_{ij}$ , then can data acquired in one sweep be leveraged in another to reduce the need for new calculations? We are interested in whether the explored parameters can span a latent space with overlapping regions that can be exploited for computational efficiency in future calculations.

(3) **Accuracy preservation:** Is the accuracy in physical quantities preserved when completing a calculation using the gGA + AL method? While it is always feasible to refine the results through a few standard gGA iterations without active learning at the end of any gGA + AL calculation, achieving high accuracy directly with AL is preferable for maximizing computational gains.

### B. Benchmarks for Hubbard model on infinite-coordination Bethe lattice

In this subsection we present benchmarks of our gGA + AL method as applied to the Hubbard model represented by the following Hamiltonian:

$$\hat{H} = \frac{U}{2} \sum_i (\hat{n}_i - 1)^2 - t \sum_{(i,j)} \sum_{\sigma=\uparrow,\downarrow} (c_{i\sigma}^\dagger c_{j\sigma} + \text{H.c.}), \quad (45)$$

where  $t$  the hopping between nearest-neighbor sites, and the hopping parameters are finite only between nearest-neighbor sites on a infinite-coordination Bethe lattice, with semicircular density of states  $\rho(\omega) = 2\sqrt{D^2 - \omega^2}/(\pi D^2)$ .

The energy measure is set in unit of the half-bandwidth  $D \propto t$ .

#### 1. Efficiency starting with empty database

For each  $\Delta U$ , we start with an empty database and execute a sweep of calculations, as defined in Sec. V A. Immediately following this, a test sweep is performed, leveraging the data acquired during the initial sweep. Once the test sweep for a specific  $\Delta U$  is completed, the database is reset to empty, and the entire process is repeated for the subsequent  $\Delta U$  values. The results of these benchmarks are summarized in Fig. 2, where each row corresponds to a different  $\Delta U$ . The figure showcases the value of the efficiency metric  $S$  for both the original and test sweeps. The left and right columns of the figure display the  $S$  values for metallic and Mott states, respectively.

A key observation is that no additional training data are required in the test sweeps for all  $\Delta U$  values, with the only exception at  $\Delta U = 0.6$ , where a few new training data are added to the database. This confirms the ability of our framework to “learn monotonically,” in the sense outlined in Sec. V A. It is also remarkable that the computational gains achieved through our gGA + AL approach are substantial even in the initial sweep, when starting from an empty database. In fact, the average of the efficiency metric  $S$  registered throughout each sweep at  $\Delta U = 0.6, 0.3, 0.15, 0.075$  is at 83%, 73%, 56%, and 40%, respectively, when considering both the Mott and metallic phases. The fact that

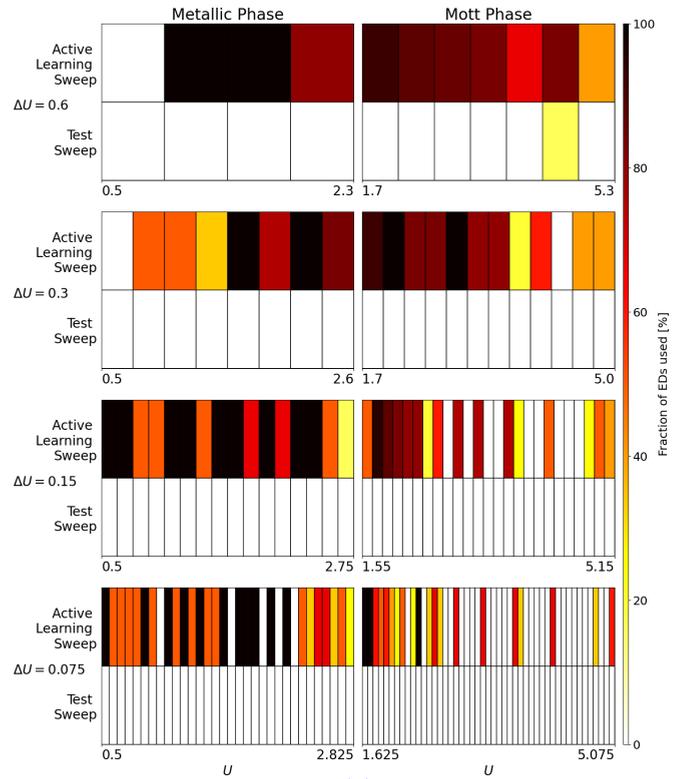


FIG. 2. Efficiency metric  $S$  for sweeps performed at different values of  $\Delta U$ . Each row corresponds to a unique  $\Delta U$ . The left and right columns show  $S$  values for metallic and Mott states, respectively. The results of both the original sweep and the subsequent test sweep are included.

tighter meshes lead to increased overall computational savings is explained by the fact that data acquired along the way for solving the gGA equations can be used by the AL framework for reducing computational cost of subsequent calculations with similar interaction strengths.

From the physical perspective, a very interesting feature of the results shown in Fig. 2 is that  $S$  is smaller in the Mott phase and the strongly correlated metallic phase, pointing to higher transfer-learning efficiency in these regions, compared to the weakly correlated regime. This finding, and its physical interpretation, is discussed later in Sec. V D with a PCA.

#### 2. Efficiency with progressive data accumulation

In Fig. 3, we compare the efficiency of our gGA + AL method across the sweeps with spacings  $\Delta U = 0.3, 0.15, 0.075$ . Each row of the figure corresponds to one of such sweeps. The top panels present the  $S$  values obtained when starting with an empty database for each new sweep, as in Fig. 2. In contrast, the bottom panels show the  $S$  values calculated starting from a database that was initially populated at the end of a  $\Delta U = 0.6$  sweep and subsequently updated without resetting as we traverse through the mentioned series of  $\Delta U$  values.

We observe that, using the scheme that retains data, the metric  $S$  demonstrates a computational gain of over 50% compared to calculations performed with a reset database.

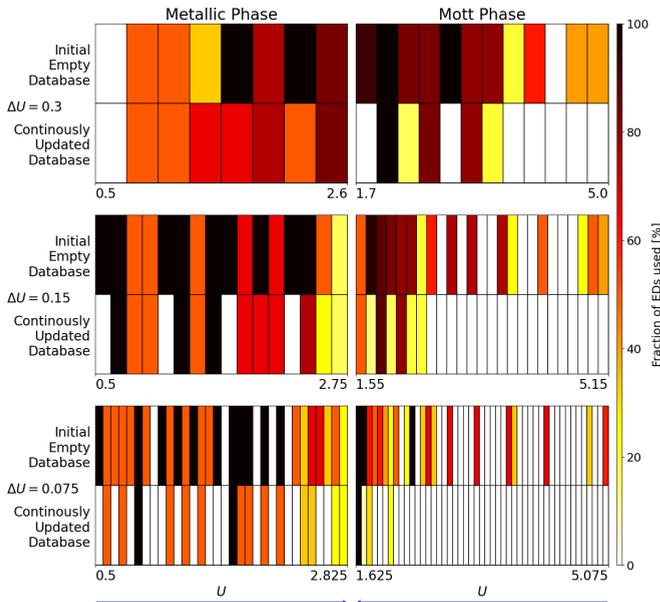


FIG. 3. Comparison of the efficiency metric  $S$  for different sweeps:  $\Delta U = 0.3, 0.15, 0.075$ . Each row corresponds to mesh spacing  $\Delta U$ . The left and right columns represent the  $S$  values for metallic and Mott states, respectively. For each  $\Delta U$ , the top panel shows  $S$  calculated using an empty database at the start of each sweep. The bottom panel displays  $S$  calculated starting from the database obtained at the end of the  $\Delta U = 0.6$  sweep and updated continuously as we proceed through the series of  $\Delta U$ , from the largest to the tightest.

Specifically, in this same scheme, the average of the efficiency metric  $S$  registers at 33%, 22%, and 10% for  $\Delta U = 0.3, 0.15, 0.075$ , respectively. Consistently with the trend of the results in Fig. 2, this gain is even more significant in the regime of strongly correlated parameters. This further supports the transfer-learning ability of our AL framework.

It is important to note that, as mentioned in Sec. III B, in the gGA framework the sequence of points  $\mathbf{X}_i$  explored throughout the computation of any given model always converges towards the point  $\bar{\mathbf{X}}$  realizing the corresponding solution. Thus, the majority of these points, where  $\mathcal{E}(\mathbf{X})$  needs to be learned, tend to cluster around  $\bar{\mathbf{X}}$  instead of being distributed throughout the entire parameter space. As a result, the data effectively form an approximate one-dimensional latent structure (parametrized by  $U$ ), embedded within the 6-dimensional ambient space, as we are going to show explicitly in Sec. V D with a PCA. The “transfer-learning” ability displayed by our AL framework (i.e., its ability to exploit the data previously stored for improving efficiency of subsequent calculations) is grounded on its ability to exploit such type of low-dimensional structures, effectively bypassing the exponential computational burden associated with the unnecessary task of learning  $\mathcal{E}(\mathbf{X})$  in the whole ambient space.

### C. Benchmarks for Hubbard model with different lattice structures

Here we extend our analysis to consider different geometries, specifically variations in the hopping matrices  $t_{ij}$ . This

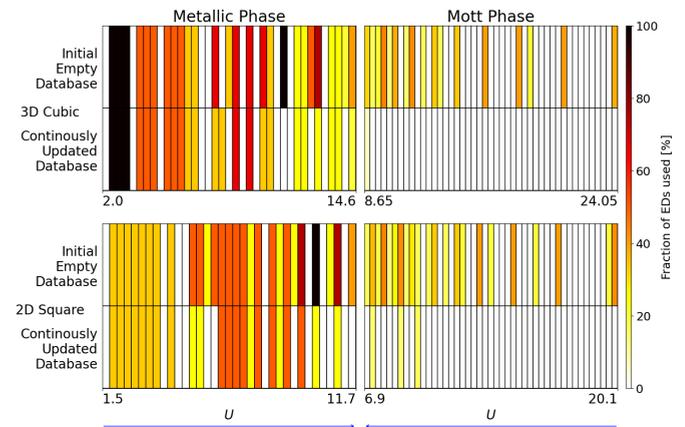


FIG. 4. Efficiency metric  $S$  for 3D cubic (upper panels) and 2D square (lower panels) lattices. Each panel is divided into two sections: The upper section shows results obtained without storing any data, while the lower section presents results where all previously acquired data, including that from the Bethe lattice calculations, were retained. Mott points are depicted on the right side, and metallic points are on the left.

enables us to evaluate how the gGA + AL framework performs when the data do not naturally span an approximately one-dimensional curve solely parametrized by  $U$ . A central question we aim to address is whether a low-dimensional structure is commonly present among the embedding parameters  $\mathbf{X}$  in these more general scenarios, and if so, whether our AL framework can leverage this structure for more efficient and accurate calculations.

#### 1. Efficiency with progressive data accumulation

To extend the scope of our analysis, we have also performed calculations of the Hubbard model on 3D cubic and 2D square lattices, employing a mesh with  $\Delta U = 0.075$ . The resulting efficiency metrics  $S$  for these calculations are illustrated in Fig. 4.

The figure is organized as follows: the upper panels correspond to the 3D cubic lattice calculations, while the lower panels are for the 2D square lattice. Within each panel, the upper and lower parts distinguish between the two modes of data acquisition. The upper part of each panel displays the efficiency metrics obtained without any stored data, whereas the lower part showcases results when retaining all previously acquired data, including that from our Bethe lattice calculations.

The results of Fig. 4 are consistent with our previous findings. Even when starting from an empty database, there is a significant computational advantage. More notably, employing the data acquisition model that continuously accumulates data results in additional gain compared to the data-reset calculations. This suggests that despite the differences in geometry between the systems, there is a degree of overlap in the data that is explored. Intriguingly, these gains are predominantly observed in the Mott phase and strongly correlated regime, implying a greater degree of overlap in these cases compared to the weakly correlated regime.

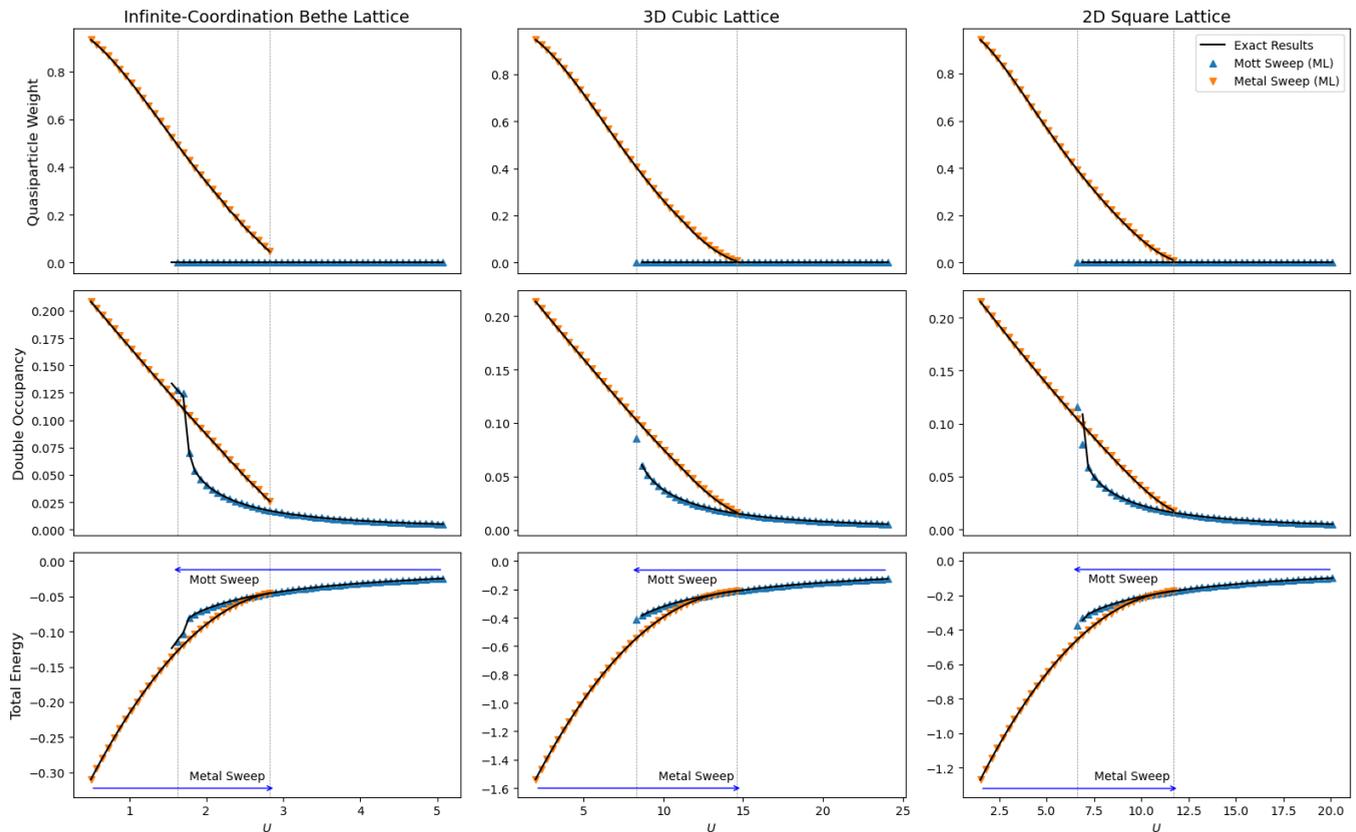


FIG. 5. Comparison of total energy, local double occupancy, and quasiparticle weight calculated using gGA + AL and standard gGA methods. The results referred to as “exact results” correspond to the exact gGA solutions and are represented by continuous black lines. The results obtained with gGA + AL are represented by triangles.

These observations lead us to conclude that the latent space of the “physically relevant embeddings,” or the set of parameters  $\mathbf{X}$  probed during gGA ground-state calculations, possesses a “special” structure. Specifically, this structure may be such that it occupies only a small subset of the ambient parameter space. We delve deeper into understanding this “special” structure of the latent space in Sec. VD, where we employ the PCA to study the data structure in detail and provide a physical interpretation of these findings.

## 2. Accuracy of gGA + AL solution

In addition to computational efficiency, another critical aspect of our gGA + AL framework is its accuracy in calculating physical observables. To rigorously evaluate this, we consider observables such as the total energy, local double occupancy, and quasiparticle weight. These observables are computed from the variational parameters obtained after convergence, as detailed in Sec. IIC.

From Fig. 5, it is evident that the application of our ML algorithm does not result in a significant loss of accuracy compared to a canonical gGA algorithm. Also, the endpoint of the metal-insulator transition  $U_{c1}$  is in perfect agreement with that obtained using the standard method. The only discrepancy is that the endpoint of the metal-insulator coexistence region  $U_{c2}$  shows a slight overestimation on the 2D square lattice when using our ML algorithm.

It is also worth pointing out that, as previously mentioned in Sec. VA, it is always possible to refine the results with

a few iterations of the standard gGA method after the active learning steps. This ensures that the computational efficiency gained by the gGA + AL framework does not compromise accuracy, providing a risk-free framework for high-efficiency, reliable calculations.

## D. PCA analysis of the training database

In this section, we turn our attention to the underlying structure of the database that has been acquired in the course of our calculations. Our primary objective is to probe the latent space within which our AL framework operates. In particular, we aim to elucidate why our AL framework shows notably higher transfer-learning efficiency in the strongly correlated regime, and to discuss the physical implications of these findings.

Our database comprises vectors  $\mathbf{X}_\alpha$  from calculations on the Bethe lattice in the limit of infinite coordination number, the 2D square lattice, and the 3D cubic lattice, all initiated without pre-existing data. To investigate the low-dimensional structure of such “latent space” of embedding parameters, we perform a PCA analysis.

### 1. Definition of the PCA

The PCA analysis of our database consists of the following steps.

(1) We construct a  $N \times d$  data matrix  $\mathbf{M}$ , where  $d = 6$  is the dimension of the vectors  $\mathbf{X}_\alpha$  and  $N$  is the number of

database points, by placing each EH parameter vector  $\mathbf{X}_\alpha$  of the database as a row, leading to:

$$\mathbf{M} = \begin{bmatrix} \vdots \\ \mathbf{X}_\alpha \\ \vdots \end{bmatrix}, \quad (46)$$

i.e.,  $\mathbf{M}_{\alpha i} = [\mathbf{X}_\alpha]_i$ .

(2) The singular value decomposition (SVD) of  $\mathbf{M}$  is represented as a sum of the outer products of its singular vectors:

$$\mathbf{M} = \sum_{r=1}^d \sigma_r \mathbf{u}_r \mathbf{v}_r^T, \quad (47)$$

where  $\sigma_r$  denotes the singular values sorted from the largest to the smallest, the vectors  $\mathbf{u}_r$  and  $\mathbf{v}_r$  are the column vectors corresponding to the left and right singular vectors, respectively. From Eq. (47) it follows that the data points  $\mathbf{X}_\alpha$  can be expressed as the following expansion of  $\mathbf{v}_r^T$ :

$$\mathbf{X}_\alpha = \sum_{r=1}^d \sigma_r [\mathbf{u}_r]_\alpha \mathbf{v}_r^T, \quad (48)$$

with coefficients  $x_{r\alpha} = \sigma_r [\mathbf{u}_r]_\alpha$ , where  $[\mathbf{u}_r]_\alpha$  denotes the  $\alpha$  component of  $\mathbf{u}_r$ . Since the singular values  $\sigma_r$  are sorted from the largest to the smallest, the first terms of Eq. (48) retain the most significant features of the parameter space.

(3) The approximation of each data point  $\mathbf{X}_\alpha$  is thus obtained by truncating the sum in Eq. (48) as follows:

$$\mathbf{X}_\alpha = \sum_{r=1}^{d_{\text{cut}}} \sigma_r [\mathbf{u}_r]_\alpha \mathbf{v}_r^T, \quad (49)$$

where  $d_{\text{cut}}$  is the number of retained principal components selected to capture the desired amount of total variance from  $\mathbf{M}$ , and  $\mathbf{v}_r^T$  are the corresponding "principal axes."

## 2. Application of the PCA

The results of the PCA analysis described above are in Fig. 6, which shows the first two principal components, i.e.,  $[\mathbf{u}_1]_\alpha \sigma_1$  and  $[\mathbf{u}_2]_\alpha \sigma_2$ , where  $\sigma_1 = 77.8$  and  $\sigma_2 = 14.6$ . These two principal components account for more than 88% of the variability of the data.

To further investigate the low-dimensional structure of the latent space, we present a scatter plot of these first two principal components in Fig. 6, providing us with a pictorial representation of the latent space. In this plot, the points are color-coded based on the lattice type and the stage of the calculation. Specifically, points obtained after convergence for the Bethe lattice, 2D square lattice, and 3D cubic lattice are colored in blue, red, and green, respectively. All other points, which are gathered during the self-consistency procedure but do not correspond to converged solutions, are colored in gray.

In line with our earlier discussion in Sec. V B 2, the data for each lattice effectively form an approximate one-dimensional latent curve, parametrized by  $U$ , which bifurcates within the coexistence region. Remarkably, data subsets corresponding to each lattice structure are very similar. Furthermore, we observe that the separation between the data corresponding to different lattices is more pronounced in the weakly correlated regime (small  $U$ , lower-left part of the graph). As  $U$  increases,

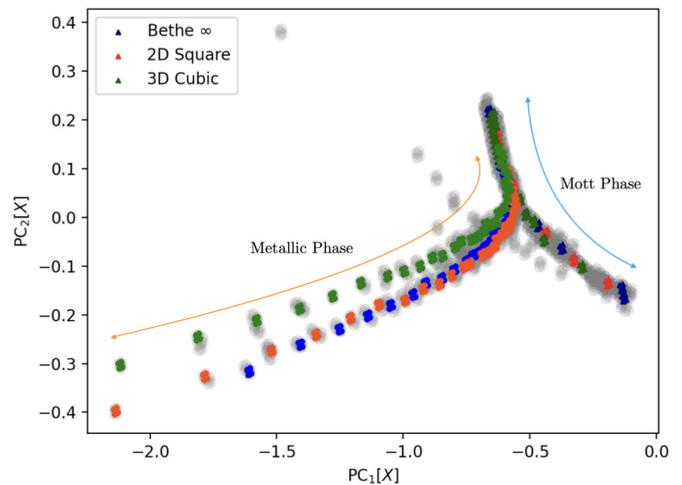


FIG. 6. Scatter plot of the first two principal components of the training database. Points obtained after convergence for the Bethe lattice, 2D square lattice, and 3D cubic lattice are colored in blue, red, and green, respectively. All other points are colored in gray.

the data corresponding to different lattices in the metallic regime become increasingly overlapping, culminating in maximum overlap near the end of the coexistence region  $U_{c2}$ .

From the computational perspective, these observations shed light on the higher transfer-learning efficiency of our AL framework in the strongly correlated regime. The overlapping data imply that similar regions of the feature space are explored across different calculations. Consequently, the data from one calculation can be effectively transferred to subsequent calculations, reducing the need for additional data points.

From a physical standpoint, the observed overlapping behavior of the databases across different lattice structures is rooted in the generality of Mott physics. The parameters  $\mathbf{X}_\alpha$  obtained after convergence, represented by the colored dots in Fig. 6, can be interpreted as physical embeddings of the correlated fragments. These embeddings capture the essence of electron localization induced by the Hubbard interaction, which transcends the specifics of the lattice structure. While in the weakly correlated regime the influence of the different lattice structures on the fragment's electronic structure is relatively pronounced, requiring to sample different regions of the ambient space, as we approach and enter the Mott phase in the strongly correlated regime, the fragments become less entangled with their surrounding environment due to reduced charge fluctuations. Consequently, it is understandable that the physical embeddings become increasingly less dependent on the lattice type, leading to the observed overlaps in the databases.

## 3. Possible future methodological enhancements

In light of our findings concerning the low-dimensional latent space, and their general origin rooted in Mott physics, it is natural to consider additional computational techniques that could further leverage this structure in future applications to complex multiorbital strongly correlated systems. Specifically, deep kernel learning with autoencoders could further facilitate learning within our AL framework, as it could offer

enhanced flexibility and scalability for discovering optimal feature spaces, all while preserving the essential element of uncertainty quantification employed within our AL procedure.

## VI. CONCLUSIONS AND OUTLOOK

In this study, we have presented an AL framework integrated within the gGA to efficiently explore the ground state of the EH in the context of the single-band Hubbard model. From a computational standpoint, this approach leads to a marked reduction in the number of EH instances that must be solved iteratively, thus significantly mitigating the computational cost inherent to gGA. Moreover, our PCA analysis reveals that the parameters of the EH reside in a latent space with a low-dimensional structure that is largely invariant to the specifics of the lattice geometry, especially in the strongly correlated regime.

From the physical perspective, the existence of this low-dimensional latent space can be attributed to the universal features of Mott physics. The phenomenon of electron localization, caused by the Hubbard interaction, transcends the geometric specifics of the systems studied. As a result, the correlated fragments are less susceptible to environmental influences, leading to a latent space whose characteristics are conserved across various lattice structures.

Looking forward, extending this methodology to more complex systems involving multiple orbitals, such as 5-orbital *d* systems and 7-orbital *f* systems, presents an interesting challenge. While the universality of Mott physics gives us reason to expect similar low-dimensional structures in these more complicated systems, the actual existence and dimensionality of such a latent space remains an open question. In this respect, it is important to note that, when considering real-material calculations, the parameters of the electronic Hamiltonian are not freely adjustable. Structural stability, which emerges from the interplay between electronic and lattice degrees of freedom, imposes further constraints on the

physically realizable electronic structures, which do not exist in model calculations, where all parameters can be tuned in arbitrary ways. A trivial example of how structural stability limits the possible quantum embeddings of the correlated degrees of freedom is that it often leads to symmetry, which can be exploited to reduce the number of gGA parameters using group-theoretical considerations. Additionally, structural stability restricts the possible atomic environments based on fundamental principles of chemistry, such as valence compatibility between atoms. These constraints may significantly limit the dimensionality and structure of the latent space of physically realizable embeddings, facilitating the learning problem.

Hence, the implementation of our AL framework in real-material calculations, potentially within an *ab-initio* DFT+gGA framework, could provide further insights into the structure of this latent space and its limitations, laying the groundwork for more efficient simulations of complex strongly correlated materials.

## ACKNOWLEDGMENTS

This work was supported by a grant from the Simons Foundation (Grant No. 1030691, N.L.). We gratefully funding from the Novo Nordisk Foundation through the Exploratory Interdisciplinary Synergy Programme Project No. NNF19OC0057790. D.G.A. acknowledges funding from the Rising Star Fellowship of the Department of Biology, Chemistry, Pharmacy of Freie Universität Berlin. The part of the work by Y.Y. was supported by the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences, Materials Science and Engineering Division, including the grant of computer time at the National Energy Research Scientific Computing Center (NERSC) in Berkeley, California. This part of research was performed at the Ames National Laboratory, which is operated for the U.S. DOE by Iowa State University under Contract No. DE-AC02-07CH11358. K.B. acknowledges support from the LANL LDRD Program.

- 
- [1] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
  - [2] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
  - [3] P. R. Kent and G. Kotliar, Toward a predictive theory of correlated materials, *Science* **361**, 348 (2018).
  - [4] Q. Sun and G. K.-L. Chan, Quantum embedding theories, *Acc. Chem. Res.* **49**, 2705 (2016).
  - [5] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* **68**, 13 (1996).
  - [6] V. Anisimov and Y. Izyumov, *Electronic Structure of Strongly Correlated Materials* (Springer, Berlin, 2010).
  - [7] K. Held, A. Nekrasov, G. Keller, V. Eyert, N. Blümer, A. K. McMahan, R. T. Scalettar, T. Pruschke, V. I. Anisimov, and D. Vollhardt, Realistic investigations of correlated electron systems with LDA+DMFT, *Phys. Status Solidi B* **243**, 2599 (2006).
  - [8] V. I. Anisimov, A. I. Oteryaev, M. A. Korotin, A. O. Anokhin, and G. Kotliar, First-principles calculations of the electronic structure and spectra of strongly correlated systems: Dynamical mean-field theory, *J. Phys.: Condens. Matter* **9**, 7359 (1997).
  - [9] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Electronic structure calculations with dynamical mean-field theory, *Rev. Mod. Phys.* **78**, 865 (2006).
  - [10] G. Knizia and G. K.-L. Chan, Density matrix embedding: A simple alternative to dynamical mean-field theory, *Phys. Rev. Lett.* **109**, 186404 (2012).
  - [11] S. Wouters, C. A. Jiménez-Hoyos, Q. Sun, and G. K.-L. Chan, A practical guide to density matrix embedding theory in quantum chemistry, *J. Chem. Theory Comput.* **12**, 2706 (2016).

- [12] F. Lechermann, A. Georges, G. Kotliar, and O. Parcollet, Rotationally invariant slave-boson formalism and momentum dependence of the quasiparticle weight, *Phys. Rev. B* **76**, 155102 (2007).
- [13] T. Li, P. Wölfle, and P. J. Hirschfeld, Spin-rotation-invariant slave-boson approach to the Hubbard model, *Phys. Rev. B* **40**, 6817 (1989).
- [14] N. Lanatà, Y. Yao, X. Deng, V. Dobrosavljević, and G. Kotliar, Slave boson theory of orbital differentiation with crystal field effects: Application to  $\text{UO}_2$ , *Phys. Rev. Lett.* **118**, 126401 (2017).
- [15] M. C. Gutzwiller, Correlation of electrons in a narrow  $s$  band, *Phys. Rev.* **137**, A1726 (1965).
- [16] W. Metzner and D. Vollhardt, Correlated lattice fermions in  $d = \infty$  dimensions, *Phys. Rev. Lett.* **62**, 324 (1989).
- [17] X.-Y. Deng, L. Wang, X. Dai, and Z. Fang, Local density approximation combined with Gutzwiller method for correlated electron systems: Formalism and applications, *Phys. Rev. B* **79**, 075114 (2009).
- [18] K. M. Ho, J. Schmalian, and C. Z. Wang, Gutzwiller density functional theory for correlated electron systems, *Phys. Rev. B* **77**, 073101 (2008).
- [19] N. Lanatà, H. U. R. Strand, X. Dai, and B. Hellsing, Efficient implementation of the Gutzwiller variational method, *Phys. Rev. B* **85**, 035133 (2012).
- [20] N. Lanatà, Y. X. Yao, C.-Z. Wang, K.-M. Ho, and G. Kotliar, Phase diagram and electronic structure of praseodymium and plutonium, *Phys. Rev. X* **5**, 011008 (2015).
- [21] J. Büneemann, F. Gebhard, T. Ohm, S. Weiser, and W. Weber, Gutzwiller-correlated wave functions: application to ferromagnetic nickel, in *Frontiers in Magnetic Materials*, edited by A. Narlikar (Springer, Berlin, 2005), p. 117.
- [22] C. Attaccalite and M. Fabrizio, Properties of Gutzwiller wave functions for multiband models, *Phys. Rev. B* **68**, 155117 (2003).
- [23] N. Lanatà, T.-H. Lee, Y.-X. Yao, and V. Dobrosavljević, Emergent Bloch excitations in Mott matter, *Phys. Rev. B* **96**, 195126 (2017).
- [24] M. S. Frank, T.-H. Lee, G. Bhattacharyya, P. Ki H. Tsang, V. L. Quito, V. Dobrosavljević, O. Christiansen, and N. Lanatà, Quantum embedding description of the Anderson lattice model with the ghost Gutzwiller approximation, *Phys. Rev. B* **104**, L081103 (2021).
- [25] C. Mejuto-Zaera and M. Fabrizio, Towards efficient computational screening of strongly correlated materials—Multiorbital phenomenology within the ghost Gutzwiller approximation, *Phys. Rev. B* **107**, 235150 (2023).
- [26] T.-H. Lee, C. Melnick, R. Adler, N. Lanatà, and G. Kotliar, Accuracy of ghost-rotationally-invariant slave-boson theory for multiorbital Hubbard models and realistic materials, *Phys. Rev. B* **108**, 245147 (2023).
- [27] T.-H. Lee, N. Lanatà, and G. Kotliar, Accuracy of ghost rotationally invariant slave-boson and dynamical mean field theory as a function of the impurity-model bath size, *Phys. Rev. B* **107**, L121104 (2023).
- [28] N. Lanatà, Derivation of the ghost Gutzwiller approximation from quantum embedding principles: Ghost density matrix embedding theory, *Phys. Rev. B* **108**, 235112 (2023).
- [29] T. Ayrál, T.-H. Lee, and G. Kotliar, Dynamical mean-field theory, density-matrix embedding theory, and rotationally invariant slave bosons: A unified perspective, *Phys. Rev. B* **96**, 235139 (2017).
- [30] T.-H. Lee, T. Ayrál, Y.-X. Yao, N. Lanatà, and G. Kotliar, Rotationally invariant slave-boson and density matrix embedding theory: Unified framework and comparative study on the one-dimensional and two-dimensional Hubbard model, *Phys. Rev. B* **99**, 115129 (2019).
- [31] L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, Machine learning for many-body physics: The case of the Anderson impurity model, *Phys. Rev. B* **90**, 155136 (2014).
- [32] E. Sheridan, C. Rhodes, F. Jamet, I. Rungger, and C. Weber, Data-driven dynamical mean-field theory: An error-correction approach to solve the quantum many-body problem using machine learning, *Phys. Rev. B* **104**, 205120 (2021).
- [33] J. Rogers, T.-H. Lee, S. Pakdel, W. Xu, V. Dobrosavljević, Y.-X. Yao, O. Christiansen, and N. Lanatà, Bypassing the computational bottleneck of quantum-embedding theories for strong electron correlations with machine learning, *Phys. Rev. Res.* **3**, 013101 (2021).
- [34] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.6.013242> for technical details about GPR extension with gradient information and the gauge fixing procedure in gGA for streamlining the ML procedure and additional benchmark calculations of the Hubbard model away from half-filling.
- [35] G.-W. Chern, K. Barros, C. D. Batista, J. D. Kress, and G. Kotliar, Mott transition in a metallic liquid: Gutzwiller molecular dynamics simulations, *Phys. Rev. Lett.* **118**, 226401 (2017).
- [36] J.-P. Julien, J. D. Kress, and J.-X. Zhu, Explicit inclusion of electronic correlation effects in molecular dynamics, *Phys. Rev. B* **96**, 035111 (2017).
- [37] H. Suwa, J. S. Smith, N. Lubbers, C. D. Batista, G.-W. Chern, and K. Barros, Machine learning for molecular dynamics with strongly correlated electrons, *Phys. Rev. B* **99**, 161107(R) (2019).
- [38] G. Schmitz, E. L. Klinting, and O. Christiansen, A Gaussian process regression adaptive density guided approach for potential energy surface construction, *J. Chem. Phys.* **153**, 064105 (2020).
- [39] D. G. Artiukhin, I. H. Godtlielsen, G. Schmitz, and O. Christiansen, Gaussian process regression adaptive density-guided approach: Towards calculations of potential energy surfaces for larger molecules, *J. Chem. Phys.* **159**, 024102 (2023).
- [40] O. Christiansen, D. G. Artiukhin, I. H. Godtlielsen, E. M. Gras, W. Györfy, M. B. Hansen, M. B. Hansen, E. L. Klinting, J. Kongsted, C. König, D. Madsen, N. K. Madsen, K. Monrad, G. Schmitz, P. Seidler, K. Sneskov, M. Sparta, B. Thomsen, D. Toffoli, and A. Zocante, MidasCpp (version 2022.10.0), <https://midascpp.gitlab.io/>.