

Surrogate models for quantum spin systems based on reduced-order modelingMichael F. Herbst^{*} and Benjamin Stamm[†]*Department of Mathematics, RWTH Aachen University, Schinkelstraße 2, 52062 Aachen, Germany*Stefan Wessel[‡]*Institute for Theoretical Solid State Physics, RWTH Aachen University, Otto-Blumenthal-Straße 26, 52074 Aachen, Germany*Matteo Rizzi[§]*Forschungszentrum Jülich GmbH, Institute of Quantum Control, Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany
and Institute for Theoretical Physics, University of Cologne, D-50937 Köln, Germany*

(Received 25 November 2021; accepted 9 March 2022; published 4 April 2022)

We present a methodology to investigate phase diagrams of quantum models based on the principle of the reduced basis method (RBM). The RBM is built from a few ground-state snapshots, i.e., lowest eigenvectors of the full system Hamiltonian computed at well-chosen points in the parameter space of interest. We put forward a greedy strategy to assemble such a small-dimensional basis, i.e., to select where to spend the numerical effort needed for the snapshots. Once the RBM is assembled, physical observables required for mapping out the phase diagram (e.g., structure factors) can be computed for any parameter value with a modest computational complexity, considerably lower than the one associated to the underlying Hilbert space dimension. We benchmark the method in two test cases, a chain of excited Rydberg atoms and a geometrically frustrated antiferromagnetic two-dimensional lattice model, and illustrate the accuracy of the approach. In particular, we find that the ground-state manifold can be approximated to sufficient accuracy with a moderate number of basis functions, which increases very mildly when the number of microscopic constituents grows—in stark contrast to the exponential growth of the Hilbert space needed to describe each of the few snapshots. A combination of the presented RBM approach with other numerical techniques circumventing even the latter big cost, e.g., tensor network methods, is a tantalizing outlook of this work.

DOI: [10.1103/PhysRevE.105.045303](https://doi.org/10.1103/PhysRevE.105.045303)**I. INTRODUCTION**

The study of quantum spin models, dating back to the early days of quantum mechanics, is a central topic in modern condensed matter physics. Indeed, as basic quantum many-body systems with inherently strong correlations, these models often display interesting ground states including complex ordering patterns, quantum disordered regimes, or topological order such as in quantum spin liquids [1]. In addition, these ground states are in many cases good approximations to the low-temperature behavior of real physical systems or compounds that are described by these models. While exact analytical solutions for specific quantum spin models exist, such as the early Bethe-ansatz solution of the one-dimensional Heisenberg spin- $\frac{1}{2}$ chain [2], most realistic models instead require advanced computational techniques for their solution.

Most traditionally, a (more or less theory-guided) scan of numerical instances needs to be computed across the parameter space of the Hamiltonian in order to map out the phase

diagram of the model—commonly through the computation of relevant observables (e.g., structure factors). Whatever is the method of choice, each such instance is typically very expensive, despite remarkable progresses to get around the naive exponential growth of the Hilbert space with the number of spins in a finite-size sample (see below). Moreover, the scan could be certainly performed in an embarrassingly parallel manner and/or initial guesses could be recycled from already converged simulations for nearby parameter values. Still, these approaches are burning a considerable amount of CPU hours. Moreover, employing previous simulations as a guess is not unproblematic and may, for example, give rise to spurious hysteresis in the proximity of phase transitions.

In this paper, we put forward an alternative and complementary strategy, relying on the so-called reduced basis method (RBM)—borrowed from the numerical mathematics community (see below)—that promises to tear down the overall amount of expensive computational instances for a faithful phase diagram. The core idea is to establish cheap surrogate models for a many-query context of parametrized quantum spin models, based on a few sample points for which snapshots, i.e., solutions of the true Hamiltonian problem, are actually produced. First, in an offline or training phase, these sample points are chosen—typically via a greedy strategy—and a reduced basis is assembled out of the corresponding

^{*}herbst@acom.rwth-aachen.de[†]best@acom.rwth-aachen.de[‡]wessel@physik.rwth-aachen.de[§]m.rizzi@fz-juelich.de

snapshots. The aim of the procedure is to obtain a subspace, which encompasses an accurate approximation of the solution across the given parameter domain. Then, in a so-called online phase, the special (affine) form of the parametrized Hamiltonians allows to obtain an approximate solution for any parameter value in a complexity solely depending on the assembled low-dimensional space and not on the potentially very high dimension of the Hilbert space.

Using such small-dimensional effective models in engineering, physics, and numerical modeling is an old idea, which has fruited not only in theoretical considerations, but also computational implementation. Early contributions in the field of such reduced-order modeling [3–6] include applications to partial differential equations (PDEs) that originate in structural and fluid mechanics. Between 2000 and 2010 the method gained much popularity due to the mathematically rigorous error control through the variational framework [7] alongside which the notion of a reduced basis method was established. Nowadays, reviews [8] as well as monographs [9,10] are available on the topic. In the context of eigenvalue problems the application of the RBM is less popular, despite the fact that the idea was already used in early contributions. See, for example, Ref. [11] for an application in structural mechanics using an effective small-dimensional basis and following a variational ansatz. For parametric eigenvalue problems in a PDE setting this idea was then formalized in Ref. [12] and extended in Ref. [13] using *a posteriori* error estimators. In Ref. [14] a generalization to target clusters of eigenvalues of a parametrized eigenvalue problem was presented.

Here, as anticipated, we want to bring the benefits of the reduced basis method to the quantum world and exploit it to speed up and economize the parameter scans required for the generation of quantum phase diagrams. As a proof of concept, we test the methodology on two quantum-spin models, namely, a model used for a chain of excited Rydberg atoms and a model for the antiferromagnetically coupled spin- $\frac{1}{2}$ triangles in the compound $\text{La}_4\text{Cu}_3\text{MoO}_{12}$. Noticeably, for a given precision target, the number of required snapshots stays moderate even when the investigated region spans different phases of the model. Furthermore, we will demonstrate this number to grow only weakly with the system size. Both features are somehow pleasantly surprising and they open up a number of theoretical questions. Moreover, the RBM strategy could contribute to green computing as well.

Before delving into the technical presentation, let us note that the RBM framework is agnostic to the precise numerical technique employed for obtaining the snapshots. It is thus fully complementary to existing algorithms. To keep this benchmark study simple we will employ exact numerical diagonalization (ED) [15,16], which is however strongly limited by the exponential growth of the Hilbert space with the number of spins in a finite-size sample. One possibility to evade such restriction is stochastic sampling of the wavefunction via quantum Monte Carlo (QMC) methods [16]; however, in the presence of geometric frustration, quantum Monte Carlo methods typically suffer severely from the negative-sign problem [17]. A modern alternative—less prone, if not immune, to such issues—is offered by the density matrix renormalization group (DMRG) [18] and its descendant tensor-network

(TN) approaches [19–22]. These are based on the insight that physically relevant states are typically low entangled and thus occupy only a comparably small subset of the total Hilbert space. Of note this is a different concept of space reduction compared to the RBM approach with the relation between the two still remaining to be explored.

The paper is organized as follows. In Sec. II, we formalize the quantum spin problems as abstract parametrized eigenvalue problems and show how the two test examples, i.e., the chain of excited Rydberg atoms and the geometrically frustrated antiferromagnetic two-dimensional lattice model, can be cast in this framework. Section III introduces the reduced basis method for the abstract family of model problems with particular focus on handling degenerate states that might appear. Section IV presents the numerical results for the two test cases while Sec. V is left for conclusions.

II. PROBLEM SETTING

We first introduce an abstract form to describe quantum spin models, and later show how our two concrete examples can be cast into this form. We consider a system with N_{sites} degrees of freedom (e.g., quantum spins), such that the quantum state of the total system belongs to the Hilbert space $\mathcal{H} = \mathbb{C}^{\mathcal{N}}$, with $\mathcal{N} = d^{N_{\text{sites}}}$, and where d denotes the dimension of the Hilbert space of each of the N_{sites} degrees of freedom. The interactions are modeled by a Hamiltonian $\mathbf{H}(\boldsymbol{\mu})$ formulated in a so-called affine decomposition, i.e.,

$$\mathbf{H}(\boldsymbol{\mu}) = \sum_{q=1}^Q \theta_q(\boldsymbol{\mu}) \mathbf{H}_q, \quad (1)$$

where $\boldsymbol{\mu} \in \mathbb{P}$ denotes a set of parameters in the parameter domain of interest, $\theta_q : \mathbb{P} \rightarrow \mathbb{R}$ indicates a scalar parameter-dependent function, and $\mathbf{H}_q : \mathcal{H} \rightarrow \mathcal{H}$ are Hermitian matrices of dimension $\mathcal{N} \times \mathcal{N}$. We assume that the number of terms, Q , is independent of \mathcal{N} .

We are now interested to evaluate, for each $\boldsymbol{\mu} \in \mathbb{P}$, the ground state(s) of the Hamiltonian $\mathbf{H}(\boldsymbol{\mu})$, i.e., the eigenvector(s) $\boldsymbol{\Psi}(\boldsymbol{\mu}) = (\Psi_1(\boldsymbol{\mu}), \dots, \Psi_m(\boldsymbol{\mu}))$ corresponding to the smallest eigenvalue $\lambda(\boldsymbol{\mu})$ solution to the problem

$$\mathbf{H}(\boldsymbol{\mu})\boldsymbol{\Psi}(\boldsymbol{\mu}) = \lambda(\boldsymbol{\mu})\boldsymbol{\Psi}(\boldsymbol{\mu}). \quad (2)$$

Note that our numerical RBM will naturally allow for (m -fold) degenerate ground states, which can arise due to symmetries.

In most applications, however, one is not so much interested in the high-dimensional solution $\boldsymbol{\Psi}(\boldsymbol{\mu})$ itself, but rather in the (scalar) expectation values of a set of physical observables as the parameters are tuned around, i.e., into a collection $\mathcal{O}(p) : \mathbb{P} \rightarrow \mathbb{R}$ of functionals of $\lambda(\boldsymbol{\mu})$ and $\boldsymbol{\Psi}(\boldsymbol{\mu})$ (p being the index running over the collection). Although a complete abstract framework is possible, we restrict ourselves to affine-decomposable observables, i.e.,

$$\mathcal{O}(p) = \sum_{r=1}^R \alpha_r(\boldsymbol{\mu}; p) \mathbf{O}_r, \quad (3)$$

where $\alpha_r(\boldsymbol{\mu}; p) \in \mathbb{C}$ are scalar coefficients and \mathbf{O}_r parameter-independent matrices of dimension $\mathcal{N} \times \mathcal{N}$. Furthermore, we

consider only averages over the ground manifold:

$$\mathcal{O}(\boldsymbol{\mu}; p) = \sum_{r=1}^R \alpha_r(\boldsymbol{\mu}; p) \frac{1}{m} \sum_{i=1}^m \Psi_i(\boldsymbol{\mu})^\dagger \mathcal{O}_r \Psi_i(\boldsymbol{\mu}). \quad (4)$$

In the applications below, for example, p will take the role of the Fourier wave vector for the structure factor as observable.

We will address the particular scenario where the solution to these eigenvalue problems needs to be solved for many different parameter values $\boldsymbol{\mu}$ in a many-query context, for example, a sweep through the parameter space, i.e., evaluating the map $\boldsymbol{\mu} \mapsto \mathcal{O}(\boldsymbol{\mu}; p)$ for many parameter values $\boldsymbol{\mu} \in \mathbb{P}$. Despite the fact that state-of-the-art numerical methods do not scale exponentially with the number of constituents—in contrast to the Hilbert-space dimension—a fine-resolution scan over the parameter domain is still very expensive due to the overall complexity to compute the ground state for *each* new parameter value. In such cases of a many-query context, we will make here use of the linear dependency of the different solutions (from one parameter value to the other) and thus render such computations considerably more affordable.

Let us now explicitly show how two physical problems of current interest in the community can be cast in this abstract framework.

A. Chains of excited Rydberg atoms

As a first example, we consider Rydberg atoms assembled into a regular lattice by means of optical tweezers, as it became customary in recent years, in a series of experiments of increasing relevance for quantum simulation purposes [23–27]. We pick up here the simplest—yet very insightful—setup of an equally spaced one-dimensional chain (with lattice spacing a), with atoms modeled as two-level systems coupled by an external laser with Rabi frequency Ω . The interplay between the level detuning Δ and the dipolar interaction strength between excited atoms gives rise to a wealth of different breaking patterns of the (discrete) translational symmetry, dictated by the Rydberg-blockade radius R_b [28–33].

We therefore consider the following Hamiltonian, scaled by $\hbar\Omega$ (and with $n_S = R_b/a$):

$$H(\boldsymbol{\mu}) := \frac{1}{2} \sum_r \sigma_r^x - \frac{\Delta}{\Omega} \sum_r \hat{n}_r + \sum_{r < r'} \left(\frac{n_S}{r' - r} \right)^6 \hat{n}_r \hat{n}_{r'}, \quad (5)$$

where the indices r, r' belong to the range $\mathcal{L} = \{1, \dots, N_x\}$. Here, σ_r^x and $\hat{n}_r = \frac{1}{2}(1 + \sigma_r^z)$ denote, respectively, the x Pauli matrix and the Rydberg excitation (single-particle) operators corresponding to the site r , according to the common convention

$$A_r := (\otimes^{r-1} \mathbb{1}) \otimes A \otimes (\otimes^{N_x-r} \mathbb{1}) \quad (6)$$

for any operator A acting on a single site.

Evidently, Hamiltonian (5) is already formulated as an affine decomposition [see Eq. (1)], with

$$\theta_1(\boldsymbol{\mu}) = 1, \quad \theta_2(\boldsymbol{\mu}) = -\mu_1 = -\frac{\Delta}{\Omega}, \quad \theta_3(\boldsymbol{\mu}) = \mu_2^6 = n_S^6,$$

the corresponding Hamiltonian components $H_{q \in \{1,2,3\}}$, and the parameter vector $\boldsymbol{\mu} = (\Delta/\Omega, n_S)$ being easily identified.

The identification of the different symmetry-breaking patterns is particularly transparent by looking at the structure factor,

$$S(\boldsymbol{\mu}; k) = \frac{1}{N_x} \sum_{r,r'} \exp(-i(r-r')k) \langle \hat{n}_r \hat{n}_{r'} \rangle_{\boldsymbol{\mu}}, \quad (7)$$

which therefore will constitute the (parametrized) output functional we are interested in. Let us notice that this is also already in the affine decomposition form assumed in Eq. (4), by introducing

$$\alpha_{r,r'}(k) = \frac{\exp(-i(r-r')k)}{N_x} \quad \text{and} \quad \mathcal{O}_{r,r'} = \hat{n}_r \hat{n}_{r'}. \quad (8)$$

B. Antiferromagnetic spin- $\frac{1}{2}$ triangles in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$

As a second example we consider a quantum spin- $\frac{1}{2}$ system that is composed out of a square lattice of triangular units (trimers), each containing three spins. This model was examined previously [34,35] as a basic model to describe the magnetism observed in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$ [36,37]. Due to an antiferromagnetic coupling of the trimer spins, this system is geometrically frustrated. As a result, it cannot be efficiently studied by QMC methods, for example, due to a severe sign problem. In earlier studies [34,35], its ground-state phase diagram, as a function of varying the intradimer couplings (as specified below), has been obtained using exact numerical diagonalization calculations, with one expensive numerical instance per point in the parameter space \mathbb{P} . Below, we will reexamine this model within our reduced-order modeling approach, which will allow us to obtain the same ground-state phase diagram through a much cheaper scan over \mathbb{P} with the surrogate model obtained in a (short) training phase.

We thus consider the following (scaled) Hamiltonian:

$$H(\boldsymbol{\mu}) = \sum_{\mathbf{r}} \left[\frac{J_1}{J_3} \mathbf{S}_{\mathbf{r},1} \cdot \mathbf{S}_{\mathbf{r},2} + \frac{J_2}{J_3} \mathbf{S}_{\mathbf{r},2} \cdot \mathbf{S}_{\mathbf{r},3} + \mathbf{S}_{\mathbf{r},3} \cdot \mathbf{S}_{\mathbf{r},1} + \frac{J'}{J_3} (\mathbf{S}_{\mathbf{r},3} \cdot \mathbf{S}_{\mathbf{r}+\mathbf{x},1} + \mathbf{S}_{\mathbf{r},2} \cdot \mathbf{S}_{\mathbf{r}+\mathbf{y},1} + \mathbf{S}_{\mathbf{r},2} \cdot \mathbf{S}_{\mathbf{r}+\mathbf{y},3}) \right], \quad (9)$$

for $\boldsymbol{\mu} = (\frac{J_1}{J_3}, \frac{J_2}{J_3}, \frac{J'}{J_3})$, where we decided to adopt J_3 as the unit of energy. It is evident that it is already cast in its affine decomposition form of Eq. (1) with

$$\theta_1(\boldsymbol{\mu}) = \mu_1 = \frac{J_1}{J_3}, \quad \theta_2(\boldsymbol{\mu}) = \mu_2 = \frac{J_2}{J_3},$$

$$\theta_3(\boldsymbol{\mu}) = 1, \quad \theta_4(\boldsymbol{\mu}) = \mu_3 = \frac{J'}{J_3},$$

and the $Q = 4$ Hamiltonian components directly readable from the above. The position vectors \mathbf{r} belong to a regular $(N_x \times N_y)$ lattice $\mathcal{L} = \mathbb{Z}^2 \cap [0, N_x) \times [0, N_y)$ with unit cell spanned by the unit vectors $\mathbf{x} = (1, 0)^\top$ and $\mathbf{y} = (0, 1)^\top$, basis elements denoted by $\alpha \in \{1, 2, 3\}$, and periodic boundary conditions imposed as $\mathbf{r} + N_x \mathbf{x} = \mathbf{r}$ and $\mathbf{r} + N_y \mathbf{y} = \mathbf{r} \forall \mathbf{r} \in \mathcal{L}$, respectively. Finally, $\mathbf{S}_{\mathbf{r},\alpha}$ denotes the vector of spin operators $\mathbf{S} = (S^x, S^y, S^z) = \frac{1}{2}(\sigma^x, \sigma^y, \sigma^z)$ acting according to Eq. (6) on the site $r = (\mathbf{r}, \alpha)$, where we converted tuples into a linear




FIG. 1. Schematic illustration of the antiferromagnetic spin- $\frac{1}{2}$ triangle model in the case $N_x = 3, N_y = 2$.

index ranging up to $3N_xN_y$, and

$$\mathbf{S}_{\mathbf{r},\alpha} \cdot \mathbf{S}_{\mathbf{r}',\alpha'} = S_{\mathbf{r},\alpha}^x S_{\mathbf{r}',\alpha'}^x + S_{\mathbf{r},\alpha}^y S_{\mathbf{r}',\alpha'}^y + S_{\mathbf{r},\alpha}^z S_{\mathbf{r}',\alpha'}^z.$$

Figure 1 provides a schematic illustration of the model and we refer to Ref. [34] for more details about the formulation.

The output functional we are interested in is again the structure factor

$$S(\boldsymbol{\mu}; \mathbf{k}) = \frac{1}{N_x N_y} \sum_{\mathbf{r}, \mathbf{r}'} \exp(-i(\mathbf{r} - \mathbf{r}') \cdot \mathbf{k}) \langle \bar{\mathbf{S}}_{\mathbf{r}} \cdot \bar{\mathbf{S}}_{\mathbf{r}'} \rangle_{\boldsymbol{\mu}}, \quad (10)$$

with the trimer total spin $\bar{\mathbf{S}}_{\mathbf{r}} = \bar{\mathbf{S}}_{\mathbf{r},1} + \bar{\mathbf{S}}_{\mathbf{r},2} + \bar{\mathbf{S}}_{\mathbf{r},3}$ (see Ref. [34]). By introducing

$$\alpha_{\mathbf{r},\mathbf{r}'}(\mathbf{k}) = \frac{\exp(-i(\mathbf{r} - \mathbf{r}') \cdot \mathbf{k})}{N_x N_y} \text{ and } \mathbf{O}_{\mathbf{r},\mathbf{r}'} = \bar{\mathbf{S}}_{\mathbf{r}} \cdot \bar{\mathbf{S}}_{\mathbf{r}'}, \quad (11)$$

we recover the affine decomposition assumed in Eq. (4).

III. METHOD

In this work, we want to make explicit use of the fact that the solution may—and for our considered examples indeed does—exhibit high linear dependency for similar parameter values. This is a working assumption that will be verified

using the upcoming greedy algorithm for each system individually. In such cases, a low-dimensional basis \mathbf{B} should exist to describe the overall solution manifold

$$\mathcal{M} = \{\Psi(\boldsymbol{\mu}) \in \mathcal{H} \mid \Psi(\boldsymbol{\mu}) \text{ g.s. of } H(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathbb{P}\} \subset \mathcal{H}$$

of the (possibly degenerate) ground states (g.s.) $\Psi(\boldsymbol{\mu})$ of the Hamiltonian $H(\boldsymbol{\mu})$ under variation of the parameters $\boldsymbol{\mu} \in \mathbb{P}$, i.e.,

$$\text{span } \mathbf{B} \simeq \mathcal{M}, \text{ with } \dim \mathbf{B} \ll \mathcal{N}.$$

We postulate that this can be the case even on parameter domains including nontrivial phase transitions. This working assumption is supported by numerical evidence in Fig. 2 for the two problems that we consider here as a benchmark. We provide the decay of the singular values σ_N of a snapshot matrix $A = [\Psi(\boldsymbol{\mu}_1) \mid \dots \mid \Psi(\boldsymbol{\mu}_{N_M})]$ over a test grid Ξ_{test} representing the best approximation error (in the ℓ^2 sense over Ξ_{test}) using N basis functions. The details of the computations are provided in the figure caption. In the case of the Rydberg chain, we observe that the number of basis functions required to approximate any element of the snapshot matrix for fixed tolerance only increases very mildly for increasing N_x . Note that, relative to the dimension 2^{N_x} of the underlying Hilbert space, the solution manifold is very low dimensional. This is illustrated in Fig. 3, where the value of N , for different tolerances, is plotted with respect to the size of the Rydberg chain with a direct comparison to the dimension $\mathcal{N} = 2^{N_x}$ of the Hilbert space (note the logarithmic scale for the y axis). For the frustrated triangular model, instead, we observe an apparent significant increase with respect to the parameters N_x, N_y . This can be explained by the fact that in a very specific region of the parameter domain we are considering—namely, for $J_1 = J_2$ and $J' \ll J_3$ —we observe degenerate ground states with degeneracies of degree $2^{N_x N_y}$, and this significantly increases the dimension of the solution manifold targeted by the RBM. Such behavior is, however, related to the model-dependent (and actually parameter-dependent) high frustration, and would affect any other approach with a similar impact. We included this parameter regime as a stress test for our RBM, while we could have circumvented it by

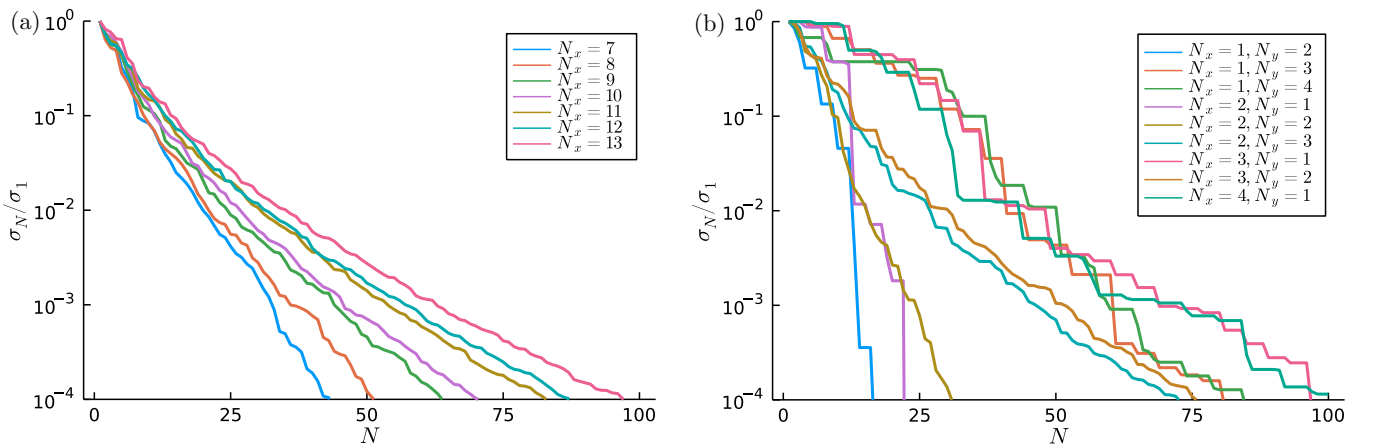


FIG. 2. Decay of normalized singular values of a snapshot matrix $A = [\Psi(\boldsymbol{\mu}_1) \mid \dots \mid \Psi(\boldsymbol{\mu}_{N_M})]$ consisting of wave functions for a given Ξ_{test} consisting of 49×49 uniformly spaced points in $\mathbb{P} = [0, 5] \times [0.5, 4]$ and $9 \times 9 \times 9$ uniformly spaced points in $\mathbb{P} = [0, 2] \times [0, 2] \times [0.01, 0.1]$ for the case of (a) Rydberg chains and (b) antiferromagnetic spin- $\frac{1}{2}$ triangles, respectively.




FIG. 3. Comparison of basis size to reach particular truncation error and size of the Hilbert space.

considering only larger values of J'/J or a training grid not including the diagonal in the J_1 - J_2 plane. Of note, even $2^{N_x N_y}$ is asymptotically fairly small compared to the dimension $2^{3N_x N_y}$ of the Hilbert space, motivating to pursue a reduced-basis approach even in this setting.

Let us be once more explicit in stating that here we deal with the complexity in the parameter space, not with the exponential cost in the physical system size. The latter can be accounted for by combining our method with a numerical method of choice in a black-box manner, whenever the numerical instance for a particular parameter value needs to be computed. Our RBM framework is therefore complementary (and not competitive) with respect to established numerical many-body methods.

Based on the postulate above, one can develop an efficient method providing an effective surrogate model for cheap scans over the entire parameter domain within a many-query context. Indeed, once such a low-dimensional approximation is constructed, the ground-state computation can simply be performed by a classical Rayleigh-Ritz procedure in this subspace. The method described here consist of an offline-online procedure as is classical in the context of the reduced-basis method [9]:

(a) *Offline.* This step consists of constructing the low-dimensional approximation space by using a greedy algorithm [9], which we recall here for the sake of completeness. Given a so-called training (or trial) grid Ξ_{train} consisting of a sizable number of M training points in \mathbb{P} , we will generate a sequence of low-dimensional approximation spaces of the form

$$\mathbb{V}_n := \text{span}\{\Psi(\mu_1), \dots, \Psi(\mu_n)\} \equiv \text{span } \mathbf{B}_n,$$

where each $\Psi(\mu_k) = (\Psi_1(\mu_k), \dots, \Psi_{m_k}(\mu_k))$, $k = 1, \dots, n$, denotes all m_k ground states of $\mathbf{H}(\mu_k)$, explicitly including the treatment of degenerate ground states in the formalism. The dimension of \mathbb{V}_n is therefore $d_n := \dim(\mathbb{V}_n) \leq \sum_{j=1}^n m_j$, and $\mathbf{B}_n \in \mathbb{C}^{N \times d_n}$ denotes a basis, for which technical details will be given later. The parameter values $\{\mu_1, \dots, \mu_n\} \subset \Xi_{\text{train}}$ (sample points) are a few well-chosen points, whose selection criterion will be described shortly as well. We call \mathbb{V}_n the reduced basis space. The construction of \mathbb{V}_n requires the com-

putation of n (truth) high-dimensional problems of the form of Eq. (2).

Starting with an initial parameter value $\mu_1 \in \Xi_{\text{train}}$, the selection of μ_k is performed by induction and is based on a greedy algorithm. Thus, let us assume that \mathbb{V}_n based on $\{\mu_1, \dots, \mu_n\}$ is given. To select μ_{n+1} we proceed as follows. We solve

$$\Phi_{\text{rb}}^{(n)}(\mu) = \arg \min_{\Phi \in \mathbb{V}_n} \frac{\Phi^\dagger \mathbf{H}(\mu) \Phi}{\Phi^\dagger \Phi} \quad (12)$$

for each $\mu \in \Xi_{\text{train}}$, where $\Phi_{\text{rb}}^{(n)}(\mu)$ denotes the vector of (possibly degenerate) ground states. Using the variational principle, the above solution can be realized by solving the d_n -dimensional (generalized) eigenvalue problem

$$\mathbf{h}(\mu) \varphi_{\text{rb}}^{(n)}(\mu) = \lambda_{\text{rb}}^{(n)}(\mu) \mathbf{b} \varphi_{\text{rb}}^{(n)}(\mu) \quad (13)$$

with $\mathbf{b} = \mathbf{B}_n^\dagger \mathbf{B}_n \in \mathbb{C}^{d_n \times d_n}$. Here $\lambda_{\text{rb}}^{(n)}(\mu)$, $\varphi_{\text{rb}}^{(n)}(\mu)$ denote explicitly the smallest eigenvalue of multiplicity m with corresponding eigenvectors $\varphi_{\text{rb}}^{(n)}(\mu) = (\varphi_{\text{rb},1}^{(n)}(\mu), \dots, \varphi_{\text{rb},m}^{(n)}(\mu))$ of $\mathbf{h}(\mu)$. The reduced (or compressed) Hamiltonian $\mathbf{h}(\mu)$ is assembled as

$$\mathbf{h}(\mu) = \sum_{q=1}^Q \theta_q(\mu) \mathbf{h}_q, \quad (14)$$

with $\mathbf{h}_q = \mathbf{B}_n^\dagger \mathbf{H}_q \mathbf{B}_n \in \mathbb{C}^{d_n \times d_n}$. The solutions to Eqs. (12) and (13) are related by the expression $\Phi_{\text{rb}}^{(n)}(\mu) = \mathbf{B}_n \varphi_{\text{rb}}^{(n)}(\mu)$. In other words $\Phi_{\text{rb}}^{(n)}(\mu)$ is the solution represented in the Hilbert space \mathcal{H} and $\varphi_{\text{rb}}^{(n)}(\mu) \in \mathbb{C}^{d_n \times m}$ collects the coefficients of all $\Phi_{\text{rb}}^{(n)}(\mu)$ when expressed in the basis \mathbf{B}_n . In accordance with the usual terminology of numerical linear algebra we will also use the terms Ritz vector and Ritz value to refer to the RBM approximations $\lambda_{\text{rb}}^{(n)}(\mu)$ and $\Phi_{\text{rb}}^{(n)}(\mu)$ of the true ground states $\lambda(\mu)$ and $\Psi(\mu)$, respectively.

To find the next sample point μ_{n+1} we select the parameter value μ for which the RBM Ritz pairs maximize the residual of the high-dimensional eigenvalue problem, i.e.,

$$\mu_{n+1} = \arg \max_{\mu \in \Xi_{\text{train}}} \text{Res}_n(\mu) \quad (15)$$

with

$$\text{Res}_n(\mu) := \sqrt{\sum_{i=1}^m \|\mathbf{H}(\mu) \mathbf{B}_n \varphi_{\text{rb},i}^{(n)}(\mu) - \lambda_{\text{rb}}^{(n)}(\mu) \mathbf{B}_n \varphi_{\text{rb},i}^{(n)}(\mu)\|_{\mathcal{H}}^2}.$$

Thanks to the *affine decomposition*, the residuals can be obtained by

$$\begin{aligned} \text{Res}_n(\mu)^2 &= \sum_{q,q'=1}^Q \theta_q(\mu)^* \theta_{q'}(\mu) \sum_{i=1}^m \varphi_{\text{rb},i}^{(n)}(\mu)^\dagger \mathbf{h}_{qq'} \varphi_{\text{rb},i}^{(n)}(\mu) \\ &\quad - \lambda_{\text{rb}}^{(n)}(\mu)^2 \sum_{i=1}^m \varphi_{\text{rb},i}^{(n)}(\mu)^\dagger \mathbf{b} \varphi_{\text{rb},i}^{(n)}(\mu). \end{aligned} \quad (16)$$

Here, the $(d_n \times d_n)$ matrices $\mathbf{h}_{qq'} := \mathbf{B}_n^\dagger \mathbf{H}_q \mathbf{H}_{q'} \mathbf{B}_n$ are parameter independent and can be computed once and for all as soon as \mathbf{B}_n is known.

Of note, once the above quantities have been computed and stored, the computational time required to assemble Eq. (14), to obtain a solution of the reduced eigenvalue problem (13),

and finally to compute the residual (16) no longer scales with \mathcal{N} , but only depends on d_n and Q . Thus, the scan of the training grid via a simple loop is of affordable complexity.

Having determined μ_{n+1} the corresponding high-dimensional eigenvalue problem (2) is solved exactly with a solver of choice. The resulting m_{n+1} ground states $\Psi(\mu_{n+1})$ are added to \mathbb{V}_n to yield the new space \mathbb{V}_{n+1} .

Let us make two remarks. First, the snapshots $\Psi(\mu_{n+1})$ might be highly linearly dependent on the existing basis \mathbf{B}_n . It is therefore advised to act upon $\Psi(\mu_{n+1})$ in order to reduce the condition number of the reduced Hamiltonian $\mathbf{h}(\mu)$, e.g., by (approximately) orthonormalizing them for defining \mathbf{B}_{n+1} . Indeed, by applying a singular value decomposition (SVD) or column-pivoted QR decomposition to the orthogonal projection onto the complement of \mathbb{V}_n one can obtain a reduced set of vectors

$$\mathbf{U}_n = \text{compress}(\Psi(\mu_{n+1}) - \mathbf{B}_n^\dagger \mathbf{b}^{-1} \mathbf{B}_n \Psi(\mu_{n+1}), \text{tol}), \quad (17)$$

in which unnecessary modes from $\Psi(\mu_{n+1})$, which differ less than the target tolerance $\text{tol} > 0$ from \mathbf{B}_n , are dropped. In the example of the SVD, this consists simply of choosing as \mathbf{U}_n those first left eigenvectors corresponding to singular values larger than tol . From \mathbf{U}_n the basis for step $n+1$ is constructed in the usual hierarchical manner concatenating $\mathbf{B}_{n+1} = [\mathbf{B}_n | \mathbf{U}_n]$.

Second, let us highlight that we do not actually need the full \mathcal{N} -dimensional formulation of the full solutions $\Psi(\mu_n)$, but merely their contractions (scalar products and expectation values) to arrive at the definition of the (compressed) matrices \mathbf{b} , \mathbf{h}_q , and $\mathbf{h}_{qq'}$. This makes the framework in principle compatible with a tensor network based solver, where $\Psi(\mu_n)$ is obtained in an economic form, which is polynomially expensive in the number of system constituents [18–22].

As already emphasized, starting with an initial parameter value $\mu_1 \in \Xi_{\text{train}}$, one repeats this procedure until the maximal residual error over the training grid is small enough. This yields an N -dimensional reduced basis based on n_t eigenvalue computations, i.e., $N = d_{n_t}$. In this manner, the high-dimensional eigenvalue problem (2) only needs to be solved n_t times. The solution of these n_t problems as well as the compression step (17) and the computation of the reduced matrices \mathbf{h}_q , $\mathbf{h}_{qq'}$, and \mathbf{b} are the only steps that depend on \mathcal{N} .

Note that the actual surrogate model, i.e., the existing reduced basis approximation $\Phi_{\text{rb}}^{(n)}(\mu) = \mathbf{B}_n \phi_{\text{rb}}^{(n)}(\mu)$ which can be cheaply computed for any μ , can be used in order to generate accurate initial guesses $\Phi_{\text{rb}}^{(n)}(\mu_{n+1})$ for the high-dimensional eigenvalue computation $\Psi(\mu_{n+1})$. Moreover, the number of points, n_t , where the greedy algorithm needs to compute the truth solution is typically much smaller than the number M of grid points in Ξ_{train} .

The algorithm is schematically summarized in Algorithm 1 where the \mathcal{N} -dependent bottlenecks are marked with a subscript \mathcal{N} .

(b) *Online*. Once the reduced basis is assembled, the so-called online step can be started. It is the evaluation of the map $\mu \mapsto \phi_{\text{rb}}(\mu)$ for any $\mu \in \mathbb{P}$ where $\phi_{\text{rb}}(\mu)$ denotes the solution to Eq. (13) and which is *independent* of \mathcal{N} (only depends on N and Q). Note that for the online step $n = N$ and we drop for the quantities in this step the superscript (N) since the reduced basis is fixed at this point.

Algorithm 1. Overview of the offline step.

Data: Training grid $\Xi_{\text{train}} \subset \mathbb{P}$, $\mu_1 \in \Xi_{\text{train}}$, n_t the number of truth eigenvalue computations.

Result: A surrogate reduced basis model $\text{rbm}_{n_t} \equiv \{\mathbf{B}_{n_t}, \mathbf{b}, \mathbf{h}_q, \mathbf{h}_{qq'}\}$ with N basis functions based on n_t truth solves.

$$\Psi(\mu_1) \leftarrow \text{truth} - \text{solver}_{\mathcal{N}}(\mu_1) \quad (2)$$

$$\text{rbm}_1 \leftarrow \text{compress}_{\mathcal{N}}(\Psi(\mu_1)) \quad (17)$$

while $\max_{\mu \in \Xi_{\text{train}}} \text{res}_n(\mu) > \text{tol}$ **do**

for $\mu \in \Xi_{\text{train}}$ **do**

$$\Phi_{\text{rb}}^{(n)}(\mu) \leftarrow \text{rb-solver}(\text{rbm}_n) \quad (13)$$

$$\text{res}_n(\mu) \leftarrow \text{residual}(\Phi_{\text{rb}}^{(n)}(\mu), \text{rbm}_n) \quad (16)$$

$$\mu_{n+1} \leftarrow \arg \max_{\mu \in \Xi_{\text{train}}} \text{res}_n(\mu) \quad (15)$$

$$\Psi(\mu_{n+1}) \leftarrow \text{truth-solver}_{\mathcal{N}}(\mu_{n+1}) \quad (2)$$

$$\mathbf{U}_n \leftarrow \text{compress}_{\mathcal{N}}(\mathbf{B}_n, \Psi(\mu_{n+1})) \quad (17)$$

$$\text{rbm}_{n+1} \equiv \{\mathbf{B}_{n+1}, \mathbf{b}, \mathbf{h}_q, \mathbf{h}_{qq'}\} \dots$$

$$\leftarrow \text{assemble}_{\mathcal{N}}(\mathbf{B}_n, \mathbf{U}_n)$$

$$\text{precompute } \mathbf{o}_r = \mathbf{B}_N^\dagger \mathbf{O}_r \mathbf{B}_N$$

From the RBM a representation of the solution in the high-dimensional Hilbert space $\mathbb{C}^{\mathcal{N}}$ can then be obtained by $\mathbf{B}_N \phi_{\text{rb}}(\mu)$, which scales linearly in \mathcal{N} . However, typically this is not needed for any practical purpose—which is one of the strengths of the RBM framework. Most importantly the computation of an output functional can indeed be performed in a complexity independent of \mathcal{N} using

$$\mathcal{O}_N(\mu; p) = \sum_{r=1}^R \alpha_r(\mu; p) \sum_{i=1}^m \phi_{\text{rb},i}(\mu)^\dagger \mathbf{o}_r \phi_{\text{rb},i}(\mu) \quad (18)$$

with precomputed $\mathbf{o}_r = \mathbf{B}_N^\dagger \mathbf{O}_r \mathbf{B}_N$. In turn, the output evaluation $\mu \mapsto \mathcal{O}_N(\mu; p)$ becomes *independent* of \mathcal{N} and only depends on N , Q , and R .

The algorithm is schematically summarized in Algorithm 2.

IV. RESULTS

We test the methodology on the two example applications introduced in Secs. II A and II B. The former is interesting as it allows to consider different chains of variable lengths and analyze the structure factors. The latter is interesting since it contains degeneracies due to symmetries allowing to test the methodology in this case.

To conduct our benchmarking tests, we implemented the method, as outlined above, in JULIA [38]. For our computation we exploit the sparsity of the Hamiltonians $\mathbf{H}(\mu)$ using compressed sparse column matrix storage. The exact ground states have been obtained iteratively using the locally optimal block preconditioned conjugate gradient

Algorithm 2. Overview of the online step.

Data: rbm_{n_t} , \mathbf{o}_r , $\mu \in \mathbb{P}$

Result: $\lambda_{\text{rb}}(\mu)$, $\mathcal{O}_N(\mu; p)$

$$\lambda_{\text{rb}}(\mu), \phi_{\text{rb}}(\mu) \leftarrow \text{rb-solver}(\text{rbm}_{n_t}) \quad (13)$$

$$\mathcal{O}_N(\mu; p) \leftarrow \text{rb-output}(\phi_{\text{rb}}(\mu)) \quad (18)$$

(LOBPCG) algorithm [39–41] as implemented in the density-functional toolkit (DFTK) [42,43] using an incomplete (sparse) Cholesky factorization as a preconditioner. Whenever a truth solve needs to be performed and a reduced basis already exists, we use the RBM approximation as a surrogate to provide an accurate initial guess. We ensure the correct number of degenerate ground states is found by adapting the number of targeted eigenpairs during the iterative diagonalization until the obtained eigenvalue gap is larger than the convergence tolerance. Our scheme has been carefully verified to obtain the correct number of eigenpairs even in the highly degenerate cases of the antiferromagnetic spin- $\frac{1}{2}$ triangles. Notice that for the chain of Rydberg atoms the LOBPCG sometimes fails to converge within a given number of iterations. In such cases we fall back to a full diagonalization of the Hamiltonian using a direct method. This happens especially frequently in the interstitial region between two clear phases. Overcoming this limitation, e.g., by including not only the ground state but also the closest manifolds of excited states in the reduced basis, is an interesting direction for future work.

In the upcoming analyses, we will use

$$\text{err}_{\text{val}} = \max_{\mu \in \Xi_{\text{test}}} \frac{|\lambda(\mu) - \lambda_{\text{rb}}(\mu)|}{|\lambda(\mu)|} \quad (19)$$

in order to assess the quality of the reduced basis approximation to the eigenvalue over a given test grid Ξ_{test} that is different from the training grid Ξ_{train} . The computation of the eigenvector error is slightly more complicated with degenerate ground states. For that purpose we will compare the spectral projectors onto the different eigenspaces given by the density matrices; i.e., we will consider the following error quantity,

$$\text{err}_{\text{vec}} = \max_{\mu \in \Xi_{\text{test}}} \frac{\|\Psi(\mu)\Psi(\mu)^\dagger - \Phi_{\text{rb}}(\mu)\Phi_{\text{rb}}(\mu)^\dagger\|_F}{\|\Psi(\mu)\Psi(\mu)^\dagger\|_F}, \quad (20)$$

where $\Phi_{\text{rb}}(\mu) = B_{n_t} \varphi_{\text{rb}}(\mu)$, the Ritz vector(s), and $\|\cdot\|_F$ denotes the Frobenius norm. Furthermore, the output functional will be measured as

$$\text{err}_{\text{sf}} = \max_{\mu \in \Xi_{\text{test}}} \frac{\|S(\mu; \cdot) - S_{\text{rb}}(\mu; \cdot)\|_F}{\|S(\mu; \cdot)\|_F}, \quad (21)$$

where the truth structure factor S and its reduced basis approximation S_{rb} are evaluated on a fine enough grid in Fourier space.

Note that we use equispaced training and test grids. As an alternative, one can use random training grids, such as proposed in Ref. [44] for greedy strategies, but one may not discover special cases and degeneracies arising at very particular parameter values, such as diagonals in parameter space, etc.

In both our test cases, we investigate the properties of the greedy algorithm, analyze the convergence of the reduced basis approximation with respect to the basis size N , and compare it with the optimal decay revealed by the singular values of the snapshot matrix discussed at the beginning of this section. Let us emphasize that the greedy algorithm does not require the truth solution at all grid points; we only use them here in order to measure the error in different metrics to assess the quality of the greedy algorithm, but this is a truly academic investigation. Furthermore, we will present




FIG. 4. Decay of the maximal residual over the training grid Ξ_{train} during the greedy algorithm with respect to N for chains of Rydberg atoms of different lengths N_x . The singular values are dotted as a reference.

results on the output functional in the form of the appropriate structure factors as defined above.

A. Chain of excited Rydberg atoms

Here, we consider a chain of excited Rydberg atoms of varying size N_x as in Eq. (5) and a parameter domain $\mathbb{P} = [0, 5] \times [0.5, 4]$, where we recall that $\mu_1 = \Delta/\Omega$ and $\mu_2 = n_S$. We apply the greedy-sampling strategy outlined in Sec. III using a training grid Ξ_{train} consisting of 50×50 uniformly spaced points in \mathbb{P} . The evolution of the maximal residual over the training grid Ξ_{train} during the offline step is illustrated in Fig. 4, where we also show the decay of the singular values. We observe that the decay of the residual nicely follows the decay of the singular values up to a constant offset. As was already observed for the decay of the singular values, the effective dimension N for fixed tolerance only increases mildly for increasing N_x .

In order to give an illustration how the greedy algorithm acts, we show in Fig. 5 the error profiles for different values of N for the one-dimensional parameter space along the line




FIG. 5. Illustration of the greedy algorithm for a one-dimensional parameter space $\{4.5\} \times [0.5, 4]$ in the case of a chain of Rydberg atoms with $N_x = 13$.




FIG. 6. Decay of the error quantities of the RBM and the residual during the greedy algorithm with respect to N for a chain of Rydberg atoms with $N_x = 13$.

of fixed $\mu_1 = 4.5$ (corresponding to the orange line in the upcoming Fig. 8). We illustrate the profile of the residual over the one-dimensional parameter space during the greedy iterations, i.e., for different values of $n = 2, 4, 6, 8$. The maximum of each curve is marked by a dot and the first eight sample points μ_n ($n = 1, \dots, 8$) by vertical lines. We observe that in agreement with theory, at each iteration n , the residuals vanish at the sample points μ_i , for $i = 1, \dots, n - 1$, since the truth solution at those points already belongs to the reduced basis space. By always adding the worst approximation (as quantified by the residual), this approach allows us to reduce the maximum residual over the parameter space and over the iterations n , though not necessarily in a monotonic fashion (see also Fig. 4).

Returning to the two-dimensional parameter domain, we test in Fig. 6 the actual errors err_{val} , err_{vec} , and err_{sf} of the eigenvalue, eigenvector, and structure factors, respectively, for an increasing sequence of N for fixed $N_x = 13$. For testing, we use a test grid Ξ_{test} consisting of 49×49 uniformly spaced points in \mathbb{P} that are different from those in the training grid Ξ_{train} . For comparison we also illustrate the evolution of the residual. We observe that the error surrogate given by the residual follows closely the real error in the eigenvector and an accelerated convergence of the eigenvalues, which is typical and expected for variational eigenvalue approximations. The fact that the eigenvector error, the residual, and the singular values exhibit the same decay rate means that the greedy strategy assembles the reduced basis in a quasioptimal fashion in this test case. Figure 7 shows the distribution of the eigenvalue error and eigenvector error over the test grid Ξ_{test} for $N = 100$ and $N_x = 13$ as well as the sample points selected by the greedy algorithm. Both eigenvalue and eigenvector errors show very similar behavior with roughly a constant offset in the error values. The larger errors in the eigenvectors are to be expected due to the faster decay in the eigenvalue error apparent in Fig. 6.

Furthermore, Fig. 8 shows the highest occupation number of the states in the canonical basis for the same parameters as considered above for $N = 20$ and $N = 100$. In the eyeball norm, one can only detect small differences, for instance in the transition between the \mathbb{Z}_2 and \mathbb{Z}_3 crystalline phases.




FIG. 7. Logarithmic distribution of the (a) eigenvalue error and (b) eigenvector error of the RBM for $N = 100$ over the parameter domain and the sample points selected by the greedy algorithm (crosses) for a chain of Rydberg atoms with $N_x = 13$. The purple dot denotes a specific parameter set that is examined in more detail in later figures.

In order to identify the dominant modes, we illustrate in Fig. 9 the structure factors along the vertical line with $\mu_1 = \Delta/\Omega = 4.5$ for $N = 100$. We observe that the RBM can capture faithfully the distinct arrangements of Rydberg excitations inside the different lobes, namely, their \mathbb{Z}_ℓ character with ℓ a divisor of $N_x - 1$. Certainly, given the number of Rydberg atoms that we consider here, one cannot expect to resolve even further details about the phase diagram such as the nature of the narrow floating phase between the lobes or the nature of the quantum phase transitions in this model [28–33].

In Fig. 10 we show the convergence of the structure factor for fixed $\mu = (4.5, 3.7)$ (corresponding to the purple point in Figs. 7 and 8 and the vertical purple line in Fig. 9) for different values of $N = 2, 4, 8$ besides the structure factors of the pure state \mathbb{Z}_4 . This point μ was chosen such that it is not close to any sampling point of the greedy assembly (see Fig. 7). We observe a very fast convergence in the eyeball norm.

B. Antiferromagnetic spin- $\frac{1}{2}$ triangles in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$

We consider the antiferromagnetic lattice model of Eq. (9) with varying sizes N_x, N_y and a parameter domain

$$\mathbb{P} = [0, 2] \times [0, 2] \times [0.01, 0.1],$$




FIG. 8. The occupation number using the RBM with (a) $N = 20$ and (b) $N = 100$ over the parameter domain for a chain of Rydberg atoms with $N_x = 13$. The orange line at $\mu_1 = 4.5$ indicates the direction along which the structure factor is plotted in Fig. 9 and the purple dot denotes the value at which its convergence is discussed in Fig. 10.

where we recall that $\mu_1 = J_1/J_3$, $\mu_2 = J_2/J_3$, and $\mu_3 = J'$. In this regime, the model constitutes an interesting stress test for our method, since it exhibits (for small J') ground-state manifolds with a degree of degeneracy that varies across the




FIG. 9. The structure factor (in logarithmic scale) using the RBM as a function of $\mu_2 = n_S$ for fixed $\mu_1 = 4.5$ and the wave number k for a chain of Rydberg atoms with $N_x = 13$. The purple line indicates the parameter value at which the convergence of the structure factor is discussed in Fig. 10.




FIG. 10. The structure factor (in logarithmic scale) using the RBM for different values of N and fixed $\mu = (4.5, 3.7)$ for a chain of Rydberg atoms with $N_x = 13$. For sake of completeness, we also illustrate the solution of the truth (exact) model and the structure factor of the pure \mathbb{Z}_4 mode. Grey crosses are used to highlight the position of the truth curve, which is almost perfectly overlaid by the result obtained in the $N = 8$ basis.

parameter domain. More specifically, we find numerically the maximal degree of degeneracy itself to increase with N_x, N_y as $2^{N_x N_y}$. Another reason why we focus on the above parameter regime is that this region was previously considered for the compound $\text{La}_4\text{Cu}_3\text{MoO}_{12}$.

Again, we apply the greedy-sampling strategy outlined in Sec. III first using a training grid Ξ_{train} consisting of $10 \times 10 \times 10$ uniformly spaced points in \mathbb{P} . The evolution of the residual during the offline step is illustrated in Fig. 11, where we also show the decay of the singular values as reference. Also in this case, we observe that the decay of the residual nicely follows the decay of the singular values. Ensured by this observation we switch to a larger training grid Ξ_{train} of $20 \times 20 \times 20$ uniformly spaced points in \mathbb{P} in the remainder of this section. Note that a ground truth computation on all 8000 parameter values is computationally very demanding for the larger triangle systems. For this reason an investigation of




FIG. 11. Decay of the residual during the greedy algorithm with respect to N for the triangle models of different lengths N_x, N_y . A uniformly spaced grid of $10 \times 10 \times 10$ points was used for training. The singular values are dotted as a reference.

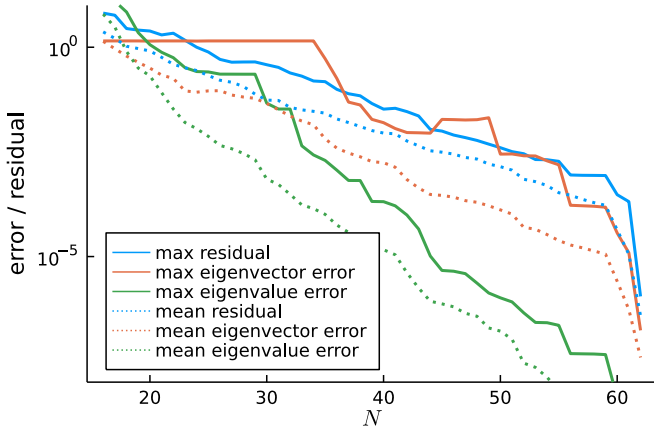


FIG. 12. Decay of the error quantities of the RBM and the residual during the greedy algorithm with respect to N for the triangular lattice model with $N_x = 2$, $N_y = 2$. A uniformly spaced grid of $20 \times 20 \times 20$ points was used for training and a testing grid of $19 \times 19 \times 19$ points distinct to Ξ_{train} . The mean values of the corresponding quantity over the test grid Ξ_{test} are dotted.

the singular value decay as well as the eigenvalue error is not attempted here.

For an RBM trained on the Ξ_{train} the decay of the approximation errors for different values of N is reported in Fig. 12 using a test grid with $19 \times 19 \times 19$ points different from Ξ_{train} . We can clearly observe a stagnation of the eigenvector error err_{vec} , measured as in Eq. (20) in the early phase of the greedy algorithm. This can be explained as follows. For small numbers N of basis functions, the approximability of the RBM is still quite inaccurate and might even fail to get the degeneracy right. As a result, the error in the density matrix $\Phi_{\text{rb}}(\mu)\Phi_{\text{rb}}(\mu)^\dagger$ is of order one. As the number N increases, this issue is cured and we observe the normal behavior of error reduction. In order to shed more light on this point, we also report in Fig. 12 the mean values of the error quantities over the test grid Ξ_{test} , i.e., replacing $\max_{\mu \in \Xi_{\text{test}}}$ by an average over Ξ_{test} , and we observe that the stagnation of the error does not occur in this mean quantity, which implies that the wrong prediction of the degeneracy only occurs within a small subset of points in Ξ_{test} .

Physically, what happens in this quantum spin model throughout most of the considered parameter space is that each triangle forms a single effective spin- $\frac{1}{2}$ degree of freedom (as long as J_1 , J_2 , and J_3 are not all equal), and these moments are then coupled by the intertriangle exchange terms $J' > 0$ to form an ordered ground state in the thermodynamic limit [34,35]. The onset of the corresponding spin ordering pattern can already be identified by examining the structure factor $\mathcal{S}(\mu; \mathbf{k})$ on comparably small system sizes. Of particular interest are the values of \mathcal{S} at the specific momenta $\mathbf{k} = (0, \pi)$ and $(\pi, 0)$, since these mixed ferromagnetic-antiferromagnetic states occupy most of the considered parameter regime μ , and we concentrate here on the results for $N_x = N_y = 2$, for which both these momenta are present for the considered periodic boundary conditions. For this model our greedy approach allows to construct small, but accurate reduced basis models. For example, with only $N = 50$ basis vectors an eigenvalue error below 10^{-6} is obtained




FIG. 13. Logarithmic distribution of the eigenvalue error of the RBM for $N = 50$ over the parameter domain and the sample points selected by the greedy algorithm (crosses) for the triangle model with $N_x = 2$, $N_y = 2$. The maximal error across all values of $\mu_3 = J'/J_3$ is shown and sampling points are projected onto the μ_1 - μ_2 plane.

across the full parameter range (compare Fig. 13). Based on an RBM with $N = 62$ we thus consider in Fig. 14 the value of the structure factor at the $\mathbf{k} = (0, \pi)$ and $(\pi, 0)$ momenta for a value of $J'/J_3 = 0.1$. We observe in particular the predominance of the $(\pi, 0)$ ordering in the upper right part of the considered parameter regime and along the line $\mu_1 = \mu_2$. An inspection of the full momentum-resolved structure factor (not shown here) indeed verifies that no other momenta provide a more dominant contribution to the magnetic structure factor. For the considered system size, the remainder of the μ_1 - μ_2 plane is dominated by the rotated $(0, \pi)$ state. We note that there is only a rather narrow region around the isotropic point $J_1 = J_2 = J_3$, for which a fully antiferromagnetic state (π, π) prevails on larger lattices [34,35].

We also tested our greedy RBM approach on larger system sizes of the triangular lattice model, namely, $N_x = 2$, $N_y = 3$ as well as $N_x = 3$, $N_y = 2$. Both settings feature a Hilbert space dimension beyond 250 000. Nevertheless, an RBM model with $N = 163$ —less than a thousandth of the Hilbert-space size—is sufficient to obtain an eigenvector residual below 0.1 and a qualitative representation of the structure factor. However, note that these system sizes are less suitable to examine which ordered phase emerges in the thermodynamic limit: systems with an odd number of unit cells in periodic boundary conditions give rise to frustration of the antiferromagnetic alignment of the trimer spins along the directions with an odd number of repeats.

Finally, let us illustrate to what extent the reduced basis is useful as an initial guess for the computation of truth solutions. Indeed, it allows us to save noticeably on the number of iterations required to obtain the converged ground states. For the triangle model with $N_x = N_y = 2$, for example, a naive guess (using a previously computed neighboring point on the grid Ξ_{test}) requires around 18 000 LOBPCG iterations to converge all ground states on our test grid of size $9^3 = 729$. In contrast employing the reduced basis with $N = 62$ as a guess one obtains all ground states using only a total number of




FIG. 14. The structure factor of the triangle model with $N_x = 2$, $N_y = 2$, and $J'/J_3 = 0.1$ at the k points (a) $(0, \pi)$ and (b) $(\pi, 0)$ for $N = 62$.

around 1600 iterations, again for all points in the grid. It turns out that the accuracy of the reduced basis is important though. For example, if only $N = 50$ is taken only 25% iterations are saved; i.e., around 13 000 iterations are needed.

V. CONCLUSIONS

In this article, we applied the reduced basis method (RBM) to parametrized quantum spin systems. Such systems naturally lead to Hamiltonians which feature an affine decomposition. In this work we exploit this structure in the context of the RBM in order to explore phase diagrams and other output quantities in a complexity that is independent of the dimension of the Hilbert space. The required reduced basis is built up using a greedy strategy based on the residual as an error surrogate, requiring only a small number of exact ground-state computations (around 20) to already reach a qualitative result. We tested the methodology for two sys-

tems, a one-dimensional chain of excited Rydberg atoms and a geometrically frustrated antiferromagnetic two-dimensional spin- $\frac{1}{2}$ lattice model.

This proof of concept was already quite conclusive. First, at the range of number of sites considered in this article, the manifold of ground states can be approximated by a surprisingly low-dimensional approximation space. From a theoretical viewpoint it would be interesting to understand the growth of the effective dimension of the solution manifold for larger numbers of particles. From a numerical perspective, the greedy strategy which uses the residual of the eigenvalue problem as error surrogate turned out to be reliable in generating reduced basis spaces in the considered tests cases. Further, the greedy algorithm was able to deal with degenerate ground states while being very efficient in the number of high-dimensional computations. Accurate reduced bases were assembled and the decay rate of the error is similar to the (optimal) one from the singular values. The reduced basis method was able to reproduce output functionals, such as structure factors or occupation numbers, fairly accurately and, e.g., only 20 basis functions (thus requiring only 20 expensive truth computations) were needed to provide a qualitatively correct occupation number plot in the case of the Rydberg chain model. Further, the method was surprisingly accurate in describing the sharp transition between the \mathbb{Z}_3 and \mathbb{Z}_4 crystalline phase with so few basis functions.

To what extent this favorable reduction of computational cost enabled by the greedy strategy generalizes to other quantum spin models is an interesting direction for future work. In particular, whether one can expect the RBM to perform well on systems, in which the degree of entanglement varies across the parameter space, is an open research problem. Our expectation is that this should not be significant with respect to the effectiveness of the RBM provided that the small corner of the Hilbert space that is actually entangled within the ground-state wave function does not vary drastically under small parameter variations. In this case a small number of truth solves should be sufficient to capture each relevant Hilbert-space corner across the parameter domain. Moreover, we emphasize that even on unseen systems the greedy strategy can be employed readily, since a careful monitoring of the decay of the residual gives direct insight whether the obtained reduced basis provides a good approximation within the studied parameter space.

One of the current limitations of our implementation, but not of the methodology itself, seems to be the use of standard eigensolvers such as the LOBPCG method to compute the (truth) ground states. As a perspective, we aim in future work to perform this task via more sophisticated tensor network methods, which efficiently deal with a larger number of microscopic constituents. They can in principle be embedded in a natural way as a black-box solver within the RBM framework, since they easily give access to the scalar products and matrix elements needed to construct the RBM surrogate model.

Another potential application of this methodology—besides the fast scan of phase diagrams—is the use of the surrogate model in order to generate initial guesses for more advanced eigenvalue solvers at a generic point in parameter space. We actually already exploited this strategy in our implementation whenever we needed to compute a truth solution

and a reduced basis was available. We demonstrated an initial guess provided by an accurate reduced basis to lead to ten times fewer iterations in the eigenvalue solver compared to the naive approach of employing the solution of a neighboring parameter value.

ACKNOWLEDGMENTS

The authors would like to thank the Center for Simulation and Data Science (CSD) of the Jülich Aachen Research Al-

liance which provided the opportunity to meet and discuss across the different disciplines. M.R. acknowledges partial support from the Deutsche Forschungsgemeinschaft (DFG), project Grant No. 277101999, within the CRC network TR 183 (Subproject No. B01), and the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 817482 (PASQuanS). S.W. acknowledges support by DFG through Grant No. WE/3649/4-2 of the FOR 1807 and through RTG 1995.

-
- [1] C. Broholm, R. J. Cava, S. A. Kivelson, D. G. Nocera, M. R. Norman, and T. Senthil, Quantum spin liquids, *Science* **367**, eaay0668 (2020).
 - [2] H. Bethe, Zur Theorie der Metalle, I. Eigenwerte und Eigenfunktionen der linearen Atomkette, *Z. Phys.* **71**, 205 (1931).
 - [3] R. Fox and H. Miura, An approximate analysis technique for design calculations, *AIAA J.* **9**, 177 (1971).
 - [4] A. K. Noor and J. M. Peters, Reduced basis technique for nonlinear analysis of structures, *AIAA J.* **18**, 455 (1980).
 - [5] A. K. Noor, On making large nonlinear problems small, *Comput. Methods Appl. Mech. Eng.* **34**, 955 (1982).
 - [6] L. Sirovich, Turbulence and the dynamics of coherent structures. I. Coherent structures, *Q. Appl. Math.* **45**, 561 (1987).
 - [7] C. Prud'homme, D. V. Rovas, K. Veroy, L. Machiels, Y. Maday, A. T. Patera, and G. Turinici, Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods, *J. Fluids Eng.* **124**, 70 (2002).
 - [8] G. Rozza, D. B. P. Huynh, and A. T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Arch. Comput. Methods Eng.* **15**, 229 (2008).
 - [9] J. S. Hesthaven, G. Rozza, and B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Vol. 590 (Springer, Berlin, 2015).
 - [10] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced Basis Methods for Partial Differential Equations: An Introduction*, Vol. 92 (Springer, Berlin, 2015).
 - [11] E. Aktas and F. Moses, Reduced basis eigenvalue solutions for damaged structures, *J. Struct. Mech.* **26**, 63 (1998).
 - [12] L. Machiels, Y. Maday, I. B. Oliveira, A. T. Patera, and D. V. Rovas, Output bounds for reduced-basis approximations of symmetric positive definite eigenvalue problems, *C. R. Acad. Sci. Ser. I Math.* **331**, 153 (2000).
 - [13] I. Fumagalli, A. Manzoni, N. Parolini, and M. Verani, Reduced basis approximation and a posteriori error estimates for parametrized elliptic eigenvalue problems, *ESAIM: Math. Modell. Numer. Anal.* **50**, 1857 (2016).
 - [14] T. Horgan, B. Wohlmuth, and T. Dickopf, Simultaneous reduced basis approximation of parameterized elliptic eigenvalue problems, *ESAIM: Math. Modell. Numer. Anal.* **51**, 443 (2017).
 - [15] A. M. Läuchli, Numerical simulations of frustrated systems, in *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, edited by C. Lacroix, P. Mendels, and F. Mila (Springer, Berlin, 2011), pp. 481–511.
 - [16] A. W. Sandvik, Computational studies of quantum spin systems, *AIP Conf. Proc.* **1297**, 135 (2010).
 - [17] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
 - [18] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
 - [19] F. Verstraete, J. I. Cirac, and V. Murg, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
 - [20] J. Eisert, Entanglement and tensor network states, in *Lecture Notes for the Autumn School on Correlated Electrons Emergent Phenomena in Correlated Matter*, edited by E. Pavarini, E. Koch, and U. Schollwöck, Modeling and Simulation Vol. 3 (Verlag des Forschungszentrum Jülich, Jülich, 2013).
 - [21] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117 (2014).
 - [22] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems, *SciPost Phys. Lect. Notes* **8** (2019).
 - [23] P. Schauß, M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch, Observation of spatially ordered structures in a two-dimensional Rydberg gas, *Nature (London)* **491**, 87 (2012).
 - [24] H. Labuhn, D. Barredo, S. Ravets, S. De Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models, *Nature (London)* **534**, 667 (2016).
 - [25] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, *Nature (London)* **551**, 579 (2017).
 - [26] P. Schauss, Quantum simulation of transverse Ising models with Rydberg atoms, *Quantum Sci. Technol.* **3**, 023501 (2021).
 - [27] M. Morgado and S. Whitlock, Quantum simulation and computing with Rydberg-interacting qubits, *AVS Quantum Sci.* **3**, 023501 (2021).
 - [28] P. Fendley, K. Sengupta, and S. Sachdev, Competing density-wave orders in a one-dimensional hard-boson model, *Phys. Rev. B* **69**, 075106 (2004).
 - [29] T. Pohl, E. Demler, and M. D. Lukin, Dynamical Crystallization in the Dipole Blockade of Ultracold Atoms, *Phys. Rev. Lett.* **104**, 043002 (2010).

- [30] M. Rader and A. M. Läuchli, Floating phases in one-dimensional Rydberg Ising chains, [arXiv:1908.02068](#).
- [31] N. Chepiga and F. Mila, Lifshitz point at commensurate melting of chains of Rydberg atoms, *Phys. Rev. Research* **3**, 023049 (2021).
- [32] N. Chepiga and F. Mila, Kibble-Zurek exponent and chiral transition of the period-4 phase of Rydberg chains, *Nat. Commun.* **12**, 414 (2021).
- [33] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, S. Choi, R. Samajdar, S. Schwartz, P. Silvi, S. Sachdev, P. Zoller, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator, *Nature (London)* **568**, 207 (2019).
- [34] S. Wessel and S. Haas, Phase diagram and thermodynamic properties of the square lattice of antiferromagnetic spin- $\frac{1}{2}$ triangles in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$, *Phys. Rev. B* **63**, 140403(R) (2001).
- [35] H.-T. Wang, Effective Hamiltonian and phase diagram of the square lattice of antiferromagnetic spin- $\frac{1}{2}$ triangles in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$, *Phys. Rev. B* **65**, 024426 (2001).
- [36] Y. Qiu, C. Broholm, S. Ishiwata, M. Azuma, M. Takano, R. Bewley, and W. J. L. Buyers, Spin-trimer antiferromagnetism in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$, *Phys. Rev. B* **71**, 214439 (2005).
- [37] M. Azuma, T. Odaka, M. Takano, D. A. Vander Griend, K. R. Poeppelmeier, Y. Narumi, K. Kindo, Y. Mizuno, and S. Maekawa, Antiferromagnetic ordering of $S = \frac{1}{2}$ triangles in $\text{La}_4\text{Cu}_3\text{MoO}_{12}$, *Phys. Rev. B* **62**, R3588 (2000).
- [38] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A fresh approach to numerical computing, *SIAM Rev.* **59**, 65 (2017).
- [39] A. V. Knyazev, Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method, *SIAM J. Sci. Comput.* **23**, 517 (2001).
- [40] U. Hetmaniuk and R. Lehoucq, Basis selection in LOBPCG, *J. Comput. Phys.* **218**, 324 (2006).
- [41] J. A. Duersch, M. Shao, C. Yang, and M. Gu, A robust and efficient implementation of LOBPCG, *SIAM J. Sci. Comput.* **40**, C655 (2018).
- [42] M. F. Herbst and A. Levitt, Density-functional toolkit (DFTK), version v0.4.0 (2021), <https://dftk.org>.
- [43] M. F. Herbst, A. Levitt, and E. Cancès, DFTK: A Julian approach for simulating electrons in solids, *Proc. JuliaCon Conf.* **3**, 69 (2021).
- [44] J. S. Hesthaven, B. Stamm, and S. Zhang, Efficient greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods, *ESAIM: Math. Modell. Numer. Anal.* **48**, 259 (2014).