Colloquium: Eigenvector continuation and projection-based emulators

Thomas Duguet

IRFU, CEA, Université Paris-Saclay, 91191 Gif-sur-Yvette, France and KU Leuven, Department of Physics and Astronomy, Instituut voor Kern- en Stralingsfysica, 3001 Leuven, Belgium

Andreas Ekström

Department of Physics, Chalmers University of Technology, SE-412 96 Göteborg, Sweden

Richard J. Furnstahl

Department of Physics, The Ohio State University, Columbus, Ohio 43210, USA

Sebastian König

Department of Physics, North Carolina State University, Raleigh, North Carolina 27695, USA

Dean Lee

Facility for Rare Isotope Beams and Department of Physics and Astronomy, Michigan State University, Michigan 48824, USA

(published 14 August 2024)

Eigenvector continuation is a computational method for parametric eigenvalue problems that uses subspace projection with a basis derived from eigenvector snapshots from different parameter sets. It is part of a broader class of subspace-projection techniques called reduced-basis methods. In this Colloquium, the development, theory, and applications of eigenvector continuation and projection-based emulators are presented. The basic concepts are introduced, the underlying theory and convergence properties are discussed, and recent applications for quantum systems and future prospects are presented.

DOI: 10.1103/RevModPhys.96.031002

CONTENTS

I. MOTIVATION

The challenges of nuclear few- and many-body physics have been addressed theoretically with a wide range of accurate but often computationally expensive *high-fidelity* methods. However, when we need to change the parameters characterizing the problem, such as Hamiltonian coupling constants, it can become computationally prohibitive to repeat high-fidelity calculations many times and challenging to reliably extrapolate. An alternative is to replace the high-fidelity model with an *emulator*, which is an approximate computer model that in the literature is sometimes referred to as a surrogate. We focus in this Colloquium on the recent development and application of emulators that exploit a technique called eigenvector continuation (EC) and its extensions. Our illustrative examples are primarily drawn from nuclear structure and reactions, for which there has been an explosion of EC applications in the last few years. We emphasize, however, the general scope of the methods, which are broadly applicable to physics problems.

Even more broadly, being able to efficiently vary the parameters in high-fidelity models to enable design, control, optimization, inference, and uncertainty quantification is a general need across engineering and science fields (Benner *et al.*, 2020a, 2020b, 2021). A common theme in these endeavors is that much of the information in high-fidelity models is superfluous. This can be exploited when tracing parametric dependencies by reducing the complexity through a so-called reduced-order model, i.e., an emulator. The universe of model order reduction methods is relatively mature but continues to expand, along with their applications.

We can put EC emulators in perspective by considering a high-level classification of reduced-order models into data-driven and model-driven categories; see Fig. 1. Data-driven methods typically interpolate the outputs of high-fidelity models without requiring an understanding of the underlying model structure; examples include Gaussian processes, artificial neural networks, and dynamic mode decomposition. Model-driven methods solve reduced-order equations derived from the full equations, so they are physics based and respect the underlying structure; examples include the broad class of reduced-basis methods (RBMs) (Hesthaven, Rozza, and Stamm, 2016; Quarteroni, Manzoni, and Negri, 2016). Increasingly there are hybrid approaches drawing from knowledge about the underlying physics problem and thereby combining both data- and model-driven aspects (Chen *et al.*, 2021).

Although originally developed independently, the model-driven EC method has long-established antecedents among RBMs [for example, eigenvalue problems in structural engineering discussed by Aktas and Moses (1998) and Nair, Keane, and Langley (1998) and in applied mathematics reviewed by Machiels *et al.* (2000), with more recent applications addressed by Fumagalli *et al.* (2016), Horger, Wohlmuth, and Dickopf (2017), and Pichi, Quaini, and Rozza (2020)]. EC uses a basis derived from selected eigenvectors from different parameter sets, called snapshots in the RBM world, to project onto a much smaller subspace than the original problem. In its simplest form, EC generates a highly effective variational basis. Typically EC applications exploit the RBM *off-line–on-line* workflow, in which expensive high-fidelity calculations are performed once in the off-line phase, enabling inexpensive but still accurate emulator calculations in the on-line phase.

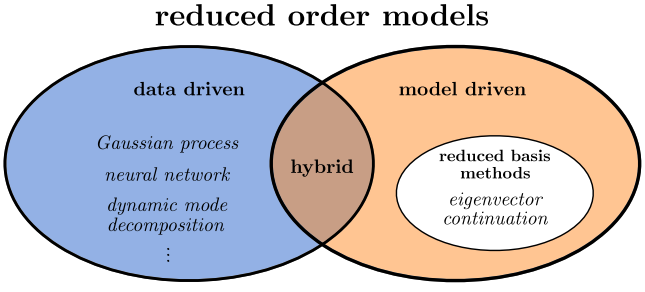

FIG. 1. Schematic classification of model order reduction emulators into data-driven methods, including Gaussian processes, artificial neural networks, and dynamic mode decomposition; model-driven methods, including reduced-basis methods (RBMs); and hybrid methods. Eigenvector continuation (EC) approaches are a subset of RBMs.

When the off-line–on-line workflow is applied to calculate observables for many parameters characterizing Hamiltonians or other operators, the EC emulators can achieve tremendous speedups over high-fidelity computational methods. This facilitates large-scale parameter exploration and calibration as well as uncertainty quantification, sensitivity analysis, and experimental design that would otherwise be infeasible. The model-driven nature of the EC approach not only ensures accurate interpolation in the parameter space but in many cases provides accurate *extrapolations* in the spaces of control parameters such as coupling strengths, energies, and boundary conditions. A consequence is that problems that are difficult or even intractable for some range of control parameters can be attacked by calculating in a range that can be more easily solved, and then extrapolating using the emulator.

Reliable emulator technology is also useful for collaboration as it enables the development of self-contained and accurate miniapplications that mimic the output of complex model calculations (Tews *et al.*, 2022; Zhang and Furnstahl, 2022). These emulators are easy to distribute given their typically small memory footprint. This allows other researchers to generate fast and accurate model predictions, even without the in-depth knowledge and significant computational resources typically required to create applications from complex or closed-source code bases.

In Sec. II we review the basic concepts of EC and the early work in nuclear physics. An overview of the RBM formulation and the off-line–on-line workflow is given in Sec. III, along with alternative approaches to generalized eigenvalue problems from a nuclear physics perspective. EC convergence properties are covered in Sec. IV, including the application to many-body perturbation theory. A major EC application is to large Hamiltonian eigensystems (Sec. V), which include adaptations to the shell model and the coupled cluster method. As illustrations of the wide range of EC applications, extensions are described for scattering (Sec. VI.A), finite-volume dependence and resonances (Sec. VI.B), and quantum Monte Carlo simulations (Sec. VI.C). A summary and a consideration of future directions are presented in Sec. VII.

II. BACKGROUND

The development of EC discussed by Frame *et al.* (2018) was inspired by the quantum many-body problem and the desire to find the extremal eigenvalues and eigenvectors of a Hamiltonian matrix too large to store in computer memory. While there are numerous quantum many-body methods used to solve such problems, they all fail when some control parameter in the Hamiltonian matrix exceeds a certain threshold value. Monte Carlo methods break down when there are strong sign oscillations and positive and negative amplitudes cancel. Simple order-by-order summations of diagrammatic expansions and perturbation theory are divergent when the magnitude of the expansion parameter exceeds unity. Variational methods are not effective if there are strong correlations not adequately captured by some wave-function ansatz or truncated set of basis states.

In the following, we review some of the concepts of EC as well as the early literature. We consider a family of matrix Hamiltonians $H(\boldsymbol{\theta})$ that depends analytically on some vector of control parameters $\boldsymbol{\theta}$, which we write in vector notation. We assume that the matrix Hamiltonians are Hermitian for all real values of the parameters. One particularly interesting and important example is the affine case where the dependence on each parameter decomposes as a sum of terms

$$H(\boldsymbol{\theta}) = \sum_{\alpha} f_{\alpha}(\boldsymbol{\theta}) H_{\alpha} \tag{1}$$

for some functions f_{α} and Hermitian matrices H_{α}. We are interested in the properties of a particular eigenvector of $H(\boldsymbol{\theta})$ and its corresponding eigenvalue $E(\boldsymbol{\theta})$,

$$H(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle = E(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle. \tag{2}$$

The basic idea of eigenvector continuation is that $|\psi(\boldsymbol{\theta})\rangle$ is an analytic function for real values $\boldsymbol{\theta}$, and the smoothness implies that it approximately lies on a linear subspace with a finite number of dimensions. We note that if there are exact eigenvalue degeneracies, the relative ordering of eigenvalues may change as we vary $\boldsymbol{\theta}$. However, the eigenvectors can still be defined as analytic functions in the neighborhood of those exact level crossings. The smoother and more gradual the undulations in the eigenvectors, the fewer the needed dimensions. A good approximation to $|\psi(\boldsymbol{\theta})\rangle$ can be efficiently found using a variational subspace composed of snapshots of $|\psi(\boldsymbol{\theta}_i)\rangle$ for parameter values $\boldsymbol{\theta}_i$. We note that for complex values of the parameters the guarantee of smoothness no longer holds.

At this point we note that other methods exist that are based on projecting a large-scale linear algebra problem into a low-dimensional subspace. Krylov methods, and, in particular, Lanczos (or, more generally, Arnoldi) iteration for calculating extremal eigenvalues of linear operators, are well established [see Saad (2011) for a textbook discussion] and used broadly, not only in physics. An important distinction compared to EC is, however, that these Krylov methods are employed at fixed $\boldsymbol{\theta}$, and thus they solve a much more limited problem. In fact, many of the EC applications discussed in Sec. V would typically use Lanczos iteration to determine the individual $|\psi(\boldsymbol{\theta})\rangle$ snapshots for the EC off-line stage.

Following Frame *et al.* (2018), we consider the Bose-Hubbard model for identical bosons on a three-dimensional cubic lattice as an illuminating and pedagogical example of EC, specializing to four bosons on a $4 \times 4 \times 4$ spatial lattice. The parameter t is the coefficient for the kinetic energy, and the parameter U is the coefficient for the pointlike interaction between pairs of bosons. For this system the relevant control parameter is the dimensionless ratio $\theta = U/t$. A variational subspace is constructed from snapshots of the eigenvectors for selected training parameters θ_j. With the shorthand $|\psi_j\rangle = |\psi(\theta_j)\rangle$, the norm matrix $\tilde{N}_{ij}$ and projected Hamiltonian matrix $\tilde{H}_{ij}(\theta)$ are given by

$$\tilde{N}_{ij} = \langle \psi_i | \psi_j \rangle, \tag{3}$$

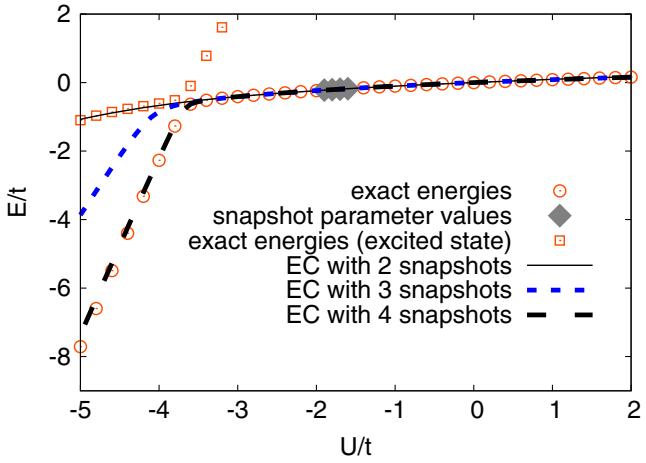

FIG. 2. Ground-state energy E of the Bose-Hubbard model divided by t vs U/t. The exact ground-state energies are shown with open circles, while the EC results are shown for variational subspace dimensions varying from 2 to 4. To highlight the avoided level crossing, the exact excited-state energies are also shown as open squares. Adapted from Frame *et al.*, 2018.

$$\tilde{H}_{ij}(\theta) = \langle \psi_i | H(\theta) | \psi_j \rangle. \tag{4}$$

The generalized eigenvalue problem is then solved as discussed in Sec. III. While "norm matrix" is the name commonly used in the nuclear physics literature for the matrix of inner products between vectors, we should note that this matrix is called the Gram matrix in the standard mathematical literature.

In Fig. 2 we show the ground-state energy E divided by t vs U/t, along with an excited state. The exact ground-state energies are shown with open circles, which reveal a sharp bend near $U/t = -3.8$. The sharp bend is caused by an avoided level crossing of eigenvalues, and the abruptness of the bend indicates that there are branch points located near the real axis. EC results are shown for subspace dimensions varying from 2 to 4. With snapshot parameter values at $U/t = -2.0, -1.9, -1.8$, and -1.7, the EC calculation is capable of extrapolating past the sharp bend.

We can understand how eigenvector continuation is able to extrapolate in this case by exploring the connection with analytic continuation. We consider a power series expansion of the eigenvector $|\psi(\theta)\rangle$ around $\theta = 0$,

$$|\psi(\theta)\rangle = \lim_{M\to\infty} \sum_{m=0}^{M} |\psi^{(m)}(0)\rangle \frac{\theta^m}{m!}. \tag{5}$$

When the series converges, we can approximate $|\psi(\theta)\rangle$ to any desired accuracy as a finite sum of $M+1$ vectors $|\psi^{(m)}(0)\rangle$, with m ranging from 0 to M.

The series will diverge when $|\theta|$ exceeds the magnitude of the nearest nonanalytic point. If $H(\theta)$ is a finite-dimensional matrix that depends analytically on θ, then the nonanalytic behavior is associated with branch points where two or more eigenvectors become the same vector (Kato, 2013). If $H(\theta)$ is a Hermitian matrix for real θ, then all of the branch points lie away from the real axis and come in complex conjugate pairs.

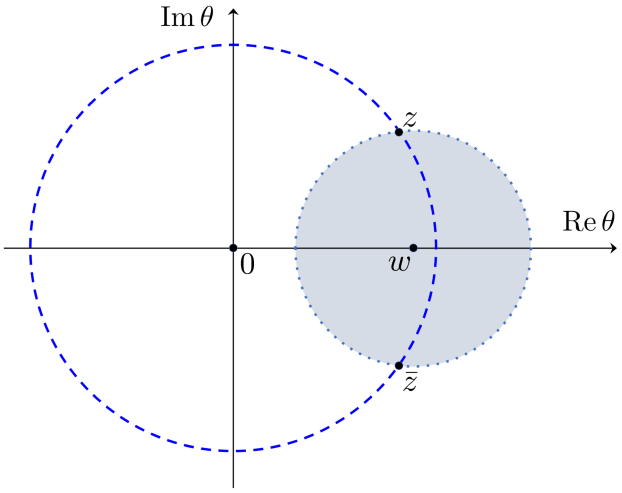

FIG. 3. While the power series expansion at $\theta = 0$ converges only for $|\theta| < |z|$, we can choose a secondary point w with $|w| < |z|$. The power series expansion at $\theta = w$ converges in the shaded region and can be reexpressed as a double series around $\theta = 0$. From Frame *et al.*, 2018.

In Fig. 3 we show an example where z and $\bar{z}$ are the branch points closest to the origin. While the power series expansion around $\theta = 0$ converges only for $|\theta| < |z|$, we can choose a secondary point w with $|w| < |z|$ and expand around w,

$$|\psi(\theta)\rangle = \lim_{N\to\infty} \sum_{n=0}^{N} |\psi^{(n)}(w)\rangle \frac{(\theta - w)^n}{n!}. \tag{6}$$

The derivatives at w can in turn be expanded using the power series about the origin. This yields the double sum

$$|\psi(\theta)\rangle = \lim_{N\to\infty} \sum_{n=0}^{N} \lim_{M\to\infty} \sum_{m=0}^{M} |\psi^{(n+m)}(0)\rangle \frac{w^m(\theta - w)^n}{m!n!}. \tag{7}$$

We can now approximate $|\psi(\theta)\rangle$ in the shaded region in Fig. 3 to any desired accuracy as a finite sum of $N + M + 1$ vectors $|\psi^{(n+m)}(0)\rangle$, with n ranging from 0 to N and m ranging from 0 to M. This process of analytic continuation shows that the approximation of $|\psi(\theta)\rangle$ using a finite linear subspace can extend beyond the nearest branch point. By including additional expansion points, this can be extended to all values of θ where $|\psi(\theta)\rangle$ is analytic.

We now return to the general problem where $H(\boldsymbol{\theta})$ depends on a vector of control parameters $\boldsymbol{\theta}$. When the eigenvector snapshots $|\psi(\boldsymbol{\theta}_i)\rangle$ are chosen with $\boldsymbol{\theta}_i$ infinitesimally close to some common limit point $\bar{\boldsymbol{\theta}}$, the variational subspace is spanned by gradients and higher-order gradients of $|\psi(\boldsymbol{\theta})\rangle$ at $\bar{\boldsymbol{\theta}}$. The EC calculation is then equivalent to a variational subspace calculation with basis states

$$\boldsymbol{\nabla}_{i_1}\boldsymbol{\nabla}_{i_2}\cdots\boldsymbol{\nabla}_{i_n}|\psi(\boldsymbol{\theta})\rangle|_{\boldsymbol{\theta}=\bar{\boldsymbol{\theta}}}. \tag{8}$$

These are the same terms that appear in the perturbation theory expansion of the eigenstate wave function (Frame *et al.*, 2018). The difference is that we are performing a variational calculation rather than evaluating partial sums of a power series. Demol *et al.* (2020) used eigenvector continuation to accelerate the convergence of Bogoliubov many-body perturbation theory; this is discussed in Sec. IV.B.

The application of EC to quantum Monte Carlo simulations was considered by Frame (2019). Since quantum Monte Carlo simulations use the Euclidean time evolution operator $e^{-H(\boldsymbol{\theta})t}$, one produces eigenstates $|\psi(\boldsymbol{\theta})\rangle$ together with exponential factors of $e^{-E(\boldsymbol{\theta})t}$. This produces some technical challenges in applications to large quantum many-body systems. The resolution of such problems is described in Sec. VI.C.

König *et al.* (2020) realized that EC could be used as a fast and accurate emulator for quantum many-body calculations by taking relatively few snapshots $\boldsymbol{\theta}_i$ to cover a compact domain of the parameter space. Ekström and Hagen (2019) extended the use of EC as an emulator to non-Hermitian matrices as encountered in coupled cluster calculations. This is discussed in Sec. V.B. Applications of EC emulators for quantum scattering problems in nuclear physics were first explored by Furnstahl *et al.* (2020). This work and several subsequent works extending the method and improving the performance are discussed in Sec. VI.A.

As noted by Bonilla *et al.* (2022) and Melendez *et al.* (2022), EC should be considered as a special case of a more general area of RBMs that has been well developed in applied mathematics over several decades, especially in the area of partial differential equations; see Hesthaven, Rozza, and Stamm (2016) and Quarteroni, Manzoni, and Negri (2016). Although the early development of EC emphasized quantum many-body systems and extrapolations for eigenvalue problems where the eigenvectors are too large to store in memory, the use of EC as an emulator is in line with many other applications of RBMs.

III. REDUCED-BASIS METHODS

The literature on RBMs is extensive (Benner *et al.*, 2020a, 2020b, 2021), with recent guides from the perspective of nuclear physicists (and EC), including pedagogical code examples, given by Melendez *et al.* (2022) and Drischler *et al.* (2023). Here we touch upon some key features that are common to EC applications.

A. RBM workflow for a Hamiltonian eigenvalue problem

The basic ingredients of an RBM workflow, which is built on a separation into off-line and on-line stages, are illustrated for a familiar Hamiltonian eigenvalue problem in Fig. 4.

Formulation in integral form. To begin, we cast the equations for the Schrödinger wave function or other quantities of interest (such as a scattering matrix) in integral form. For the Hamiltonian eigenvalue problem in Fig. 4(a) (left panel) with parameters $\boldsymbol{\theta}$, solving the finite matrix problem of a large basis size $N_h \times N_h$ is formally equivalent to finding N_h approximate stationary solutions to the variational functional

$$\mathcal{E}[\psi] = \langle\psi|H(\boldsymbol{\theta})|\psi\rangle - E(\boldsymbol{\theta})(\langle\psi|\psi\rangle - 1) \tag{9}$$

in the space spanned by the N_h basis elements. This is our high-fidelity model. (In practice, the use of previously

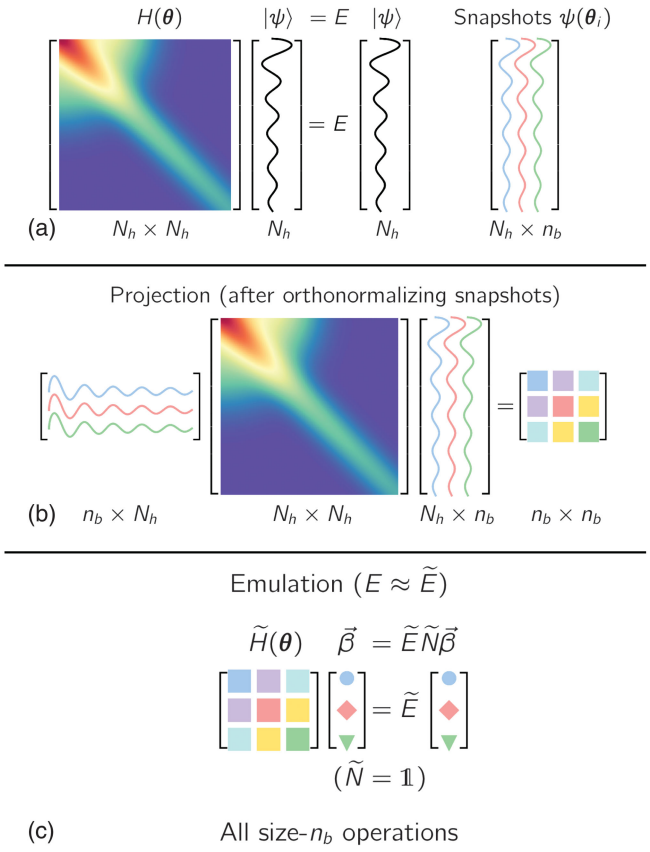

FIG. 4. Reduced-basis model workflow for a matrix eigenvalue problem. (a) High-fidelity calculations of snapshots, each with a large size N_h, (b) are projected in the off-line stage to a reduced-basis matrix of a small size $n_b \times n_b$. (c) In the on-line stage, the emulator only uses size n_b operations. Adapted from Drischler *et al.*, 2023.

mentioned Krylov methods means that finding all N_h solutions is not needed, but the computational cost still scales with N_h.) Other RBM formulations are discussed in Sec. III.B.

Off-line stage. Next we reduce the dimensionality of the problem by substituting for the general solution a trial basis of a size n_b. RBMs start with a snapshot basis consisting of high-fidelity solutions $|\psi_i\rangle$ at selected values $\{\boldsymbol{\theta}_i; i = 1, \ldots, n_b\}$ in the parameter space, as in Fig. 4(a) (right panel). When seeking the ground-state energy and wave function for arbitrary $\boldsymbol{\theta}$, these $|\psi_i\rangle$ are ground-state eigenvectors from diagonalizing $H(\boldsymbol{\theta}_i)$. For many EC applications in nuclear physics to date it has been sufficient to choose this basis randomly, for example, with a space-filling sampling algorithm such as Latin hypercube sampling. This basis spans a reduced space and can be used directly (after orthonormalizing the snapshots),

$$|\tilde{\psi}\rangle = \sum_{i=1}^{n_b} \beta_i |\psi_i\rangle, \tag{10}$$

with basis expansion coefficients β_i. The Hamiltonian is then projected onto a much smaller $n_b \times n_b$ space, as shown in Fig. 4(b).

More generally in RBM applications one first compresses the snapshot basis by applying some variant of principle component analysis [known as proper orthogonal decomposition (POD) in this context], which builds on the singular value decomposition (SVD) of the snapshots and enables a smaller basis size than n_b. Alternatively, or in conjunction with POD, one can efficiently select snapshots by applying an active learning protocol (greedy algorithm) that aims to minimize the overall error of the emulator (Hesthaven, Rozza, and Stamm, 2016). For a recent application of a greedy algorithm to quantum spin systems, including an efficient mapping of ground-state phase diagrams; see Herbst *et al.* (2022) and Brehmer *et al.* (2023). These approaches are the standard in RBM applications but are not yet widely applied in nuclear physics (Bonilla *et al.*, 2022; Sarkar and Lee, 2022).

On-line stage. For variational formulations, we enforce stationarity with respect to the trial basis expansion coefficients. This leads to an $n_b \times n_b$ generalized eigenvalue problem for the basis coefficients

$$\begin{aligned} \tilde{H}(\boldsymbol{\theta})\vec{\beta}(\boldsymbol{\theta}) &= \tilde{E}(\boldsymbol{\theta})\tilde{N}(\boldsymbol{\theta})\vec{\beta}(\boldsymbol{\theta}), \\ \tilde{H}_{ij}(\boldsymbol{\theta}) &= \langle \psi_i | H(\boldsymbol{\theta}) | \psi_j \rangle, \\ \tilde{N}_{ij}(\boldsymbol{\theta}) &= \langle \psi_i | \psi_j \rangle, \end{aligned} \tag{11}$$

as introduced in Sec. II and visualized in Fig. 4(c). Note that if the basis has been orthonormalized, then $\tilde{N}$ is an identity matrix. Extending such an emulator to matrix elements of other operators and even transitions is straightforward; see Wesolowski *et al.* (2021) for a nuclear example.

In uncertainty quantification, for which sampling of many parameter sets are usually required, it is essential that the emulator be many times faster than the high-fidelity calculations. This is achieved for RBM emulators via the off-line and on-line separation because the on-line stage requires computations scaling only with n_b (small) and not with N_h (large). An affine operator structure, meaning a factorization of parameter dependence as in Eq. (1), is needed to achieve the desired on-line efficiency because size-N_h operations such as $\langle \psi_i | H_\alpha | \psi_j \rangle$ are independent of $\boldsymbol{\theta}$ and need to be calculated only once in the off-line stage. If the problem is nonaffine, then the strategy is to apply a so-called hyper-reduction approach, which leads to an approximate affine form (Quarteroni, Manzoni, and Negri, 2016). A nuclear scattering example that treats a nonaffine Hamiltonian was given by Odell *et al.* (2024).

B. Variational and Galerkin formulations

More generally an RBM can be formulated in terms of a functional that is stationary at the desired solution (the variational approach) or via a weak form arising from multiplying the underlying equations and boundary conditions by arbitrary test functions and integrating over the relevant domains (the Galerkin approach) (Brenner and Scott, 2008; Zienkiewicz, Taylor, and Zhu, 2013; Hesthaven, Rozza, and Stamm, 2016). To date EC in nuclear physics has most often been implemented with a variational formulation, which for bound energy eigenstates is familiar from introductory physics. Less familiar but well established are various variational

approaches to quantum scattering, where each approach leads to a different emulator; see Sec. VI.A.

The Galerkin approach starts with the schematic form

$$\langle\zeta|H(\boldsymbol{\theta})-E(\boldsymbol{\theta})|\psi\rangle=0,\quad \forall\ \langle\zeta|, \tag{12}$$

with arbitrary test functions $|\zeta\rangle$. The reduced dimensionality for Galerkin RBM formulations enforces orthogonality with a restricted set of n_b test functions,

$$\langle\zeta_i|H(\boldsymbol{\theta})-\tilde{E}(\boldsymbol{\theta})|\tilde{\psi}\rangle=0,\quad i=1,\ldots,n_b. \tag{13}$$

If the test functions are chosen to be the trial basis functions $\langle\zeta_i|=\langle\psi_i|$, then they are called Bubnov-Galerkin or Ritz-Galerkin (or just Galerkin) formulations. If a different basis of test functions is used, it is called a Petrov-Galerkin formulation. For eigenvalue problems with Hamiltonians that are bounded from below, the Ritz-Galerkin procedure yields the same equations as the variational approach. The Petrov-Galerkin option means that the Galerkin procedure is more general.

For boundary-value partial or ordinary differential equations, there are general variational and Galerkin RBM formulations. A projection-based emulator seeks the solution ψ to

$$D(\psi;\boldsymbol{\theta})=0\text{ in }\Omega,\qquad B(\psi;\boldsymbol{\theta})=0\text{ on }\Gamma, \tag{14}$$

where D and B are operators in the domain Ω and its boundary Γ, respectively. There are many good references on Galerkin methods; see Brenner and Scott (2008) and Zienkiewicz, Taylor, and Zhu (2013). The canonical example of a Poisson equation with Neumann boundary conditions was worked out by Melendez *et al.* (2022).

The same RBM ingredients as in the eigenvalue problem apply here, with an integral formulation using a stationary functional such as an action $S[\psi]$, with $\delta S=0$ yielding Eq. (14), or starting with

$$\int_\Omega d\Omega\,\zeta D(\psi)+\int_\Gamma d\Gamma\bar{\zeta}B(\psi)=0, \tag{15}$$

and integrating by parts to reach a Galerkin weak formulation (Zienkiewicz, Taylor, and Zhu, 2013), which for arbitrary test functions ζ and $\bar{\zeta}$ also yields Eq. (14). With the snapshot trial basis (10), $\delta S=0$ can be solved for the optimal β_i (for linear operators this is simply a matrix equation). The Galerkin formulation needs a choice for the test basis $\langle\zeta|=\sum_{i=1}^{n_b}\delta\beta_i\langle\zeta_i|$, where $\delta\beta_i$ organize the orthogonalization conditions for each i,

$$\delta\beta_i\left[\int_\Omega d\Omega\zeta_i D(\tilde{\psi})+\int_\Gamma d\Gamma\bar{\zeta}_i B(\tilde{\psi})\right]=0. \tag{16}$$

[For notational simplicity we leave the partial integrations implicit in Eq. (16).] Again we have Ritz-Galerkin and Petrov-Galerkin options. For a broad set of engineering and science applications using these approaches, good starting points are works by Hesthaven, Rozza, and Stamm (2016), Quarteroni, Manzoni, and Negri (2016), and Benner *et al.* (2017, 2020a, 2020b). Galerkin methods for quantum scattering are discussed in Sec. VI.A, including an application to a nonaffine Hamiltonian parametrization.

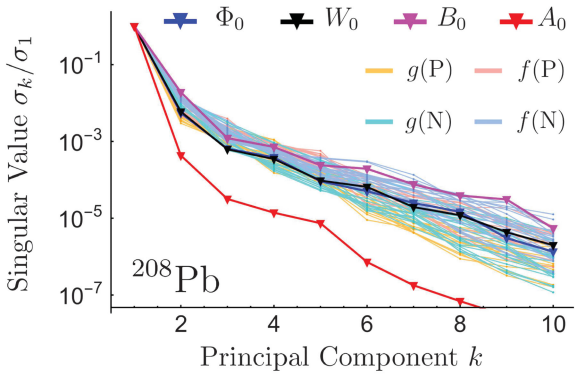

FIG. 5. Normalized singular values from $n_b=50$ snapshots of various functions that enter a nuclear physics energy density functional. The snapshots correspond to different parameter sets to be used in a Galerkin formulation of the energy density functional, which will be solved many times with different sets for Bayesian parameter estimation. The rapid decrease with principal component number k indicates that a small basis size will be accurate, leading in this case to speedups of several thousand compared to the original solver. Adapted from Giuliani *et al.*, 2023.

A pedagogical illustration of Galerkin methods adapted to nuclear physics energy density functionals for uncertainty quantification was given by Giuliani *et al.* (2023). Figure 5 shows the singular values from snapshots of various functions that arise in the coupled nonlinear differential equations to be solved in minimizing such an energy density functional. The efficacy of a basis obtained by POD from snapshots is implied by the rapid decrease in magnitude of the singular values, leading to high accuracy from a relatively small basis and speedups of the order of several thousand (the actual speedup will be implementation dependent).

C. Other approaches to generalized eigenvalue problems

As emphasized, the key to fast emulation with EC is a decomposition into one-time off-line tasks and repeated, computationally efficient on-line tasks. In this context, the key equations to be solved on-line are the low-dimensional generalized eigenvalue problem of Eq. (11) and Fig. 4(c). The nonorthogonality of the high-fidelity snapshots yields a nontrivial norm matrix $\tilde{N}(\boldsymbol{\theta})$ on the right-hand side of Eq. (11). This type of secular equation is routinely encountered in nuclear physics in the discretized version of the projected generator coordinate method (Griffin and Wheeler, 1957; Brink and Weiguny, 1968; Ring and Schuck, 1980; Frosini, Duguet, Ebran, and Somà, 2022) and in the Monte Carlo shell model (Otsuka *et al.*, 2001; Shimizu *et al.*, 2012). The same is true for the nonorthogonal configuration interaction (Thom and Head-Gordon, 2009) in quantum chemistry.

Because the norm matrix is Hermitian and the EC on-line problem is low dimensional, the norm matrix $\tilde{N}(\boldsymbol{\theta})$ can be diagonalized to transform Eq. (11) into a standard matrix

diagonalization problem.[1] Near-linear redundancies between the high-fidelity snapshots make the norm matrix poorly conditioned numerically. Consequently, its kernel L_0 must be explicitly separated from its orthogonal complement $L_\perp$ before transforming Eq. (11) unitarily. In fact, as 0 is an accumulation point of the eigenspectrum of the norm matrix in the limit of infinite dimension, small nonzero eigenvalues can generate instabilities even for the finite dimensions presently under consideration. The practical remedy to this problem consists of removing eigenvectors in L_0 associated with eigenvalues smaller than a given threshold ϵ_{th} chosen such that the end results do not depend on its specific value.[2]

There are several other methods used to deal with the inversion of poorly conditioned norm matrices. Tikhonov regularization is one popular approach (Tikhonov, 1943). The simplest form of Tikhonov regularization is ridge regression or nugget regularization. In this approach, a small positive multiple of the identity is added to the norm matrix that needs to be inverted. However, it is often not clear how to estimate the systematic bias introduced using this approach.

A new approach called the trimmed sampling algorithm was introduced by Hicks and Lee (2023). Trimmed sampling uses physics-based constraints and Bayesian inference to reduce errors of the generalized eigenvalue problem. Instead of simply regulating the norm matrix, probability distributions are sampled for the Hamiltonian and norm matrix elements, weighted by likelihood functions derived from physics-informed constraints. These physics-informed constraints include well-motivated physics principles such as positivity of the norm matrix and the smooth convergence of extremal eigenvalues with respect to variational subspace size. The posterior distribution is determined for the Hamiltonian and norm matrix elements, and eigenvectors and observables are then sampled from that distribution.

In Fig. 6 a schematic diagram of trimmed sampling is displayed. The raw uncertainty of some observable obtained from solving the generalized eigenvalue problem is sketched. The raw uncertainty centered around the starting estimate is used as the prior probability distribution. The posterior probability is proportional to the product of the prior probability and the likelihood associated with the enforcement of some physics-informed constraints. The posterior probability distribution does not give a rigorous estimate of the error. However, it can be concluded that the exact solution is located at a point where neither the prior probability nor the likelihood is small (Hicks and Lee, 2023).

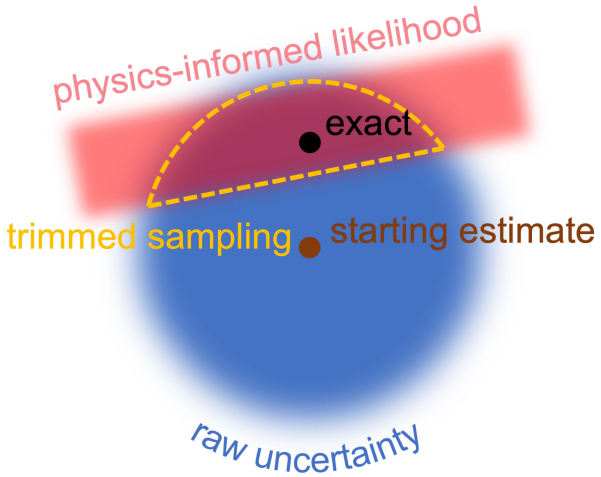

FIG. 6. Schematic of trimmed sampling. The raw uncertainty of some observable obtained from solving the generalized eigenvalue problem is sketched. The raw uncertainty is centered around the starting estimate and is used as the prior probability distribution. The posterior probability is proportional to the product of the prior probability and the likelihood associated with the enforcement of some physics-informed constraints. The exact solution is located at a point where neither the prior probability nor the likelihood is small. From Hicks and Lee, 2023.

IV. CONVERGENCE PROPERTIES OF EC

An important and fundamental question regarding EC is how fast it converges to the exact answer as a function of the number of snapshots.

A. Bounds on the EC convergence rate

We are interested in the rate of convergence of EC for interpolation as well as the more difficult problem of extrapolation. We start with the problem of interpolation and consider a Hamiltonian $H(\theta)$ with a single control parameter θ. Let B be a compact real-valued domain for θ, and let $|\psi(\theta)\rangle$ be the eigenvector of interest. Let $d(\theta, S_{n_b})$ be the norm of the residual vector when approximating $|\psi(\theta)\rangle$ using EC with snapshots $S_{n_b} = \{|\psi(\theta_1)\rangle, \ldots, |\psi(\theta_{n_b})\rangle\}$ chosen from B. We let d_{n_b} denote the best possible uniform error bound for $d(\theta, S_{n_b})$. This means that we optimize the selection of snapshots for fixed n_b such that $\max_{\theta \in B} d(\theta, S_{n_b})$ is minimized. Our d_{n_b} is an example of a Kolmogorov N width, which is used to characterize the error and convergence of linear subspace approximations (Kolmogoroff, 1936; Tikhomirov, 1960; Pinkus, 1985).

Within its radius of convergence, a power series expansion converges exponentially fast with respect to truncation order. For example, if we truncate the power series around $\theta = 0$ in Eq. (5) at order M, the resulting error will be $O(|\theta/z|^{M+1})$ in the limit of large M, where z is the nearest nonanalytic point. If $|\psi(\theta)\rangle$ is analytic on B, we can use this fundamental property of power series to derive an upper bound on EC errors when they are used for interpolation.

To begin, we select a set of points such that all parts of B lie within the radius of convergence of one of these points, which

[1]Different numerical methods are called for when the dimension is large enough to forbid a straight diagonalization of the norm matrix; see Frosini, Duguet, Ebran, Bally *et al.* (2022). Note further that the equivalence of the original and transformed secular equations is not guaranteed in the continuous version of the projected generator coordinate method (de Toledo Piza *et al.*, 1977; Broeckhove and Deumens, 1979). However, neither of these issues occurs in the present context.

[2]This is effectively equivalent to removing the singular values below ϵ_{th} in a truncated POD or SVD algorithm and is standard practice in the RBM approach.

we call anchor points. We then take snapshots at these anchor points as well as points infinitesimally close to the anchor points. Linear combinations of the snapshots can be used to construct derivatives and higher-order derivatives of $|\psi(\theta)\rangle$ at each anchor point. If we take n_b snapshots, then, in the limit of large n_b, we have enough basis vectors to express the power series at each anchor point up to a truncation order that scales as $O(n_b)$. It follows that d_{n_b} is $O(x^{n_b})$ for some positive $x < 1$. The generalization to compact real-valued domains in d dimensions is straightforward. We have $O(k^d)$ gradients and higher-order gradients for the multiparameter power series in $\theta_1, \ldots, \theta_d$ at truncation order k. For the multidimensional case, d_{n_b} is $O(x^{n_b{}^{1/d}})$ for some positive $x < 1$. In the RBM literature, snapshots composed of high-fidelity solutions and their first derivatives at different anchor points have been used for partial differential equations and this is known as the Hermite subspace approach (Ito and Ravindran, 2001).

In general, EC extrapolation converges more slowly than interpolation. Consider the case where the EC snapshots lie in the neighborhood of some point, but extrapolation is required beyond the radius of convergence at that point. As illustrated in Fig. 3, we can perform secondary expansions and analytically continue past branch points in the complex plane. We can bound the EC extrapolation error in terms of the convergence of multiseries expansions such as that shown in Eq. (7). These secondary expansions result in slower convergence, and the problem is most severe when there are branch points at $\theta = z$ and $\bar{z}$ that pinch the real axis. The number of secondary expansions needed will scale inversely with the imaginary part of z. The smaller that $\mathrm{Im}\, z$ is, the sharper the resulting avoided level crossing. For systems undergoing a quantum phase transition, $\mathrm{Im}\, z$ will decrease with system size, and this limits the ability of EC to extrapolate across the transition point in large systems (Franzke *et al.*, 2024).

The previously described analysis based on power series and perturbation theory gives an upper bound on the asymptotic error of EC for Hamiltonians $H(\boldsymbol{\theta})$ that are analytic in $\boldsymbol{\theta}$. However, the actual convergence rate of EC is typically faster than that of perturbation theory when snapshots are selected infinitesimally close to some anchor point. This stems from the fact that the gradients and higher-order gradients in Eq. (8) are not orthogonal to each other. As described by Sarkar and Lee (2021) and illustrated in Fig. 7, this results in a phenomenon that they called differential folding, where cancellations occur between terms in the power series expansion for $|\psi(\boldsymbol{\theta})\rangle$. No such phenomenon occurs in subspace-projection methods such as EC. As new snapshots are included, the linear subspace is expanding along directions that are orthogonal to the previous snapshots. This produces a faster convergence for EC than perturbation theory.

The faster convergence of EC versus perturbation theory can be seen in three different matrix examples denoted as models 1A, 1B, and 1C by Sarkar and Lee (2021). Each of these matrix models has an affine dependence on one parameter. In Fig. 8, we show the logarithm of the error for the eigenstate wave function versus truncation order for perturbation theory (PT), vector continuation (VC), and EC. Vector continuation corresponds to simple projection of the exact eigenvector onto the subspace spanned by the EC snapshots. For all three cases, we see that the VC and EC are converging significantly faster than in perturbation theory. Sarkar and Lee (2021) proved that VC and EC approximations agree up to terms that scale quadratically with the error of the VC and EC approximations.

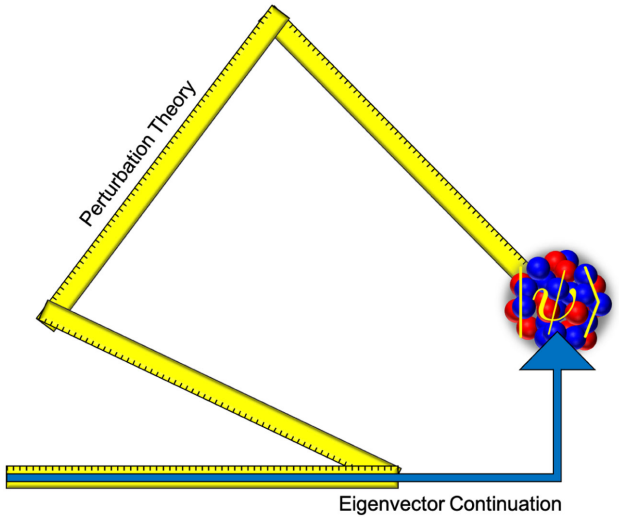

FIG. 7. The convergence of perturbation theory is impacted by a phenomenon called differential folding, where cancellations occur between terms in the power series expansion for $|\psi(\boldsymbol{\theta})\rangle$. Differential folding does not occur in EC calculations, since the linear subspace is expanding along new orthogonal directions. From Sarkar and Lee, 2021.

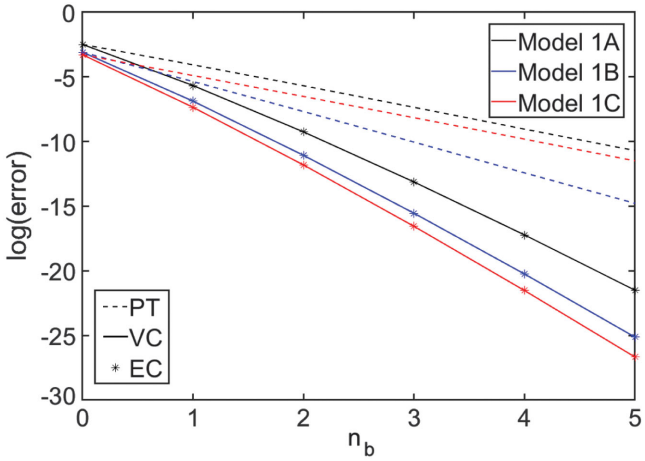

FIG. 8. The logarithm of the error for the eigenstate wave function vs truncation order for perturbation theory (PT), vector continuation (VC), and eigenvector continuation (EC). The results are for three different matrix examples described by Sarkar and Lee (2021) and labeled as models 1A, 1B, and 1C. We see that EC and VC both converge faster than in perturbation theory.

B. Improved many-body perturbation theory

A natural application of EC relates to overcoming some critical limitations of many-body perturbation theory(-ies) applied to nuclear systems. While more advanced (for

example, nonperturbative) expansion methods are currently employed to obtain accurate solutions of the nuclear many-body Schrödinger equation (Hergert, 2020), many-body perturbation theories of various flavors are of great use in many applications (Tichai, Roth, and Duguet, 2020).

In this context, the generic parametric dependence of the Hamiltonian takes the simple form

$$H(\theta) = H_0 + \theta H_1, \tag{17}$$

with θ a complex number, knowing that the case of physical interest corresponds to $\theta = 1$. Eigenstates of $H(\theta)$ can be accessed via perturbation theory as a Taylor series around $\theta = 0$, i.e., via an expansion with respect to eigenstates of H_0. Eventually an eigenstate $|\Psi_n(\theta)\rangle$ of $H(\theta)$ and its eigenenergy $E_n(\theta)$ are approximated at perturbative order P through

$$|\Psi_n^{(P)}(\theta)\rangle \equiv \sum_{p=0}^{P} \theta^p |\Phi_n^{(p)}\rangle, \tag{18}$$

$$E_n^{(P)}(\theta) \equiv \sum_{p=0}^{P} \theta^p \mathcal{E}_n^{(p)}, \tag{19}$$

where the corrections $\{(|\Phi_n^{(p)}\rangle, \mathcal{E}_n^{(p)}); p \in \mathbb{N}\}$ can be computed from the eigenstates of H_0 (Shavitt and Bartlett, 2009).

The key problem relates to the fact that the sequence $\{(|\Psi_n^{(P)}(\theta)\rangle, E_n^{(P)}(\theta)); P \in \mathbb{N}\}$ typically converges toward $(|\Psi_n(\theta)\rangle, E_n(\theta))$ when $P \to \infty$ only for $||\theta|| \in [0, R_c]$, where R_c denotes the convergence radius. If $R_c < 1$, the problem of physical interest is inaccessible via perturbation theory.

In nuclear many-body calculations, several features can lead to $R_c < 1$ (Tichai, Roth, and Duguet, 2020), such as characteristics of the internucleon interactions, the choice of H_0, and the closed- or open-shell nature of the nucleus under study. While appropriately acting on the first two aspects allows one to bypass the problem in closed-shell nuclei (Tichai *et al.*, 2016), it is much more challenging to do so in open-shell systems (Demol *et al.*, 2021). In this context, EC was shown to provide a systematic framework to enlarge the convergence radius via analytic continuation; i.e., the EC employing the set of P-order perturbative snapshots $\{|\Psi_n^{(P)}(\theta_i)\rangle; i = 1, \ldots, P+1\}$ acts as a resummation technique delivering a controlled and variational sequence of approximations to $E_n(1)$ for increasing $n_b = P + 1 \in \mathbb{N}$ (Demol *et al.*, 2020; Franzke *et al.*, 2022).

Figure 9 demonstrates that the sequence of approximations to the ground-state energy of the open-shell ^{18}O nucleus obtained via EC converges rapidly from above toward $E_0(1)$ even though the corresponding perturbative series diverges. Because the Hilbert space dimension employed is small enough, the results can be validated against the exact value of $E_0(1)$ obtained via the exact diagonalization of the nuclear Hamiltonian based on configuration interaction (CI) techniques. While the use of EC to resum diverging perturbative series was first dedicated to nuclear ground states (Demol *et al.*, 2020), it was later extended to excited eigenstates (Franzke *et al.*, 2022).

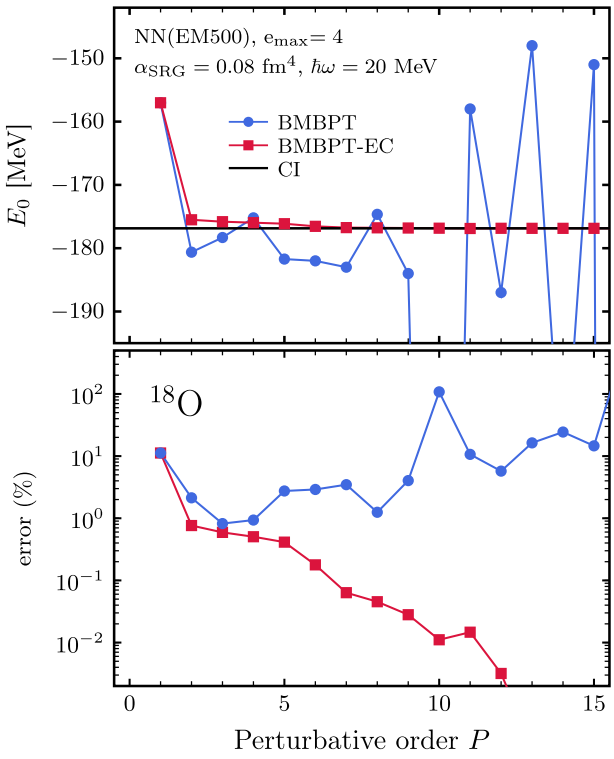

FIG. 9. Ground-state energy of ^{18}O from Bogoliubov many-body perturbation theory (BMBPT) (blue circles) and BMBPT-based EC (red squares) as a function of the perturbative order P against exact diagonalization (solid line). The employed Hilbert space dimension is small enough for the exact diagonalization of the nuclear Hamiltonian to be accessible via configuration interaction (CI) techniques. Top panel: absolute values. Bottom panel: relative error to exact diagonalization. Adapted from Demol *et al.*, 2020.

Another successful application of EC is to pairing in many-body systems; see Baran and Nichita (2023) and Franzke *et al.* (2024). However, Franzke *et al.* also manifested an aforementioned limitation of EC, as they found that they could not extrapolate between the normal and superfluid regimes if they included snapshots from only one regime. That is, extrapolating between different phases of large systems will generally fail unless information on both is included; see also Sowiński and Garcia-March (2022) and Brehmer *et al.* (2023).

V. LARGE HAMILTONIAN EIGENSYSTEMS

A powerful approach to obtaining (part of) the spectrum of the Hamiltonian of a quantum system is the explicit diagonalization of a large (typically sparse) Hamiltonian matrix. Such approaches are ideal candidates for the straightforward application of EC formulated by Frame *et al.* (2018), i.e., a Galerkin projection. Since they play a crucial role in nuclear structure theory, many related applications of EC arose relatively quickly in this context.

A. No-core shell-model emulators

As a first application that fueled many subsequent developments, König *et al.* (2020) used a no-core shell-model

framework (formulated in terms of Jacobi coordinates) to construct EC-based emulators for $A = 3$ and 4 nucleons, i.e., the nuclei ^{3}H and ^{4}He. In this approach, the wave function of the Hamiltonian, written as $H = H(\boldsymbol{\theta})$ with a collection of parameters $\boldsymbol{\theta}$, is expanded in eigenfunctions of a harmonic-oscillator potential with a chosen frequency. Truncating the harmonic-oscillator basis based on a maximum number of oscillator quanta $N_{\max}$ yields a large finite matrix that can be diagonalized. For $A = 3$ and 4 Hamiltonians formulated in Jacobi coordinates, thereby exactly factorizing the center of mass components of the wave functions, typical matrix dimensions in the large space are $N_h \times N_h = 10^4 \times 10^4$. EC for one or more states can be set up directly using the coefficient vectors obtained from Lanczos diagonalization.

The parameters $\boldsymbol{\theta}$ considered by König *et al.* (2020) are the low-energy constants of the chiral effective field theory (χEFT) potential used in that work. Overall, there are $d = 16$ individual parameters subsumed in $\boldsymbol{\theta}$ that determine two- and three-nucleon interactions in the potential. Setting up an EC emulator proceeds following the on-line–off-line scheme described in Sec. II for the generic RBM workflow.

(i) Picking a training set $\{\boldsymbol{\theta}_i\}_{i=1}^{n_b}$ of n_b parameters, using space-filling Latin hypercube sampling (McKay, Beckman, and Conover, 1979) in the d-dimensional parameter domain (or some subset thereof).
(ii) Performing exact calculations (for the ground states of ^{3}H and ^{4}He) in the case of König *et al.* (2020) for each point in the training set.
(iii) Constructing a pair of Hamiltonian and norm matrices as described in Sec. II for each evaluation of the emulator at a target parameter point $\boldsymbol{\theta}_*$.

An important property of the chiral Hamiltonian is that it typically can be written as an affine combination, as introduced in Eq. (1),

$$H(\boldsymbol{\theta}) = H_0 + \sum_{k=1}^{d} \theta_k H_k, \tag{20}$$

where H_0 denotes the kinetic energy plus parameter-independent parts of the chiral Hamiltonian. This form makes it particularly efficient to evaluate the emulator at different target points in the parameter space [step (iii) in the previous list] because each operator H_k can be individually projected onto the EC space, and this is a one-time cost that is part of the off-line emulator setup.

König *et al.* (2020) provided an analysis of the numerical performance gain (speedup factor) that is achieved via EC, shown in Fig. 10 as a particular example. While the details of that analysis are particular to the employed Jacobi-coordinate no-core shell-model calculation of ^{4}He, which is a light-mass nucleus of manageable computational complexity, much of the discussion applies generally to EC-based emulators with affine parameter dependence. Note that EC can be used to greatly reduce the effective dimension of a matrix problem, and the maximum speedup factor that follows is primarily determined by the size of the reduced dimension compared to the original one. The speedup factor as a function of the number of on-line emulator samples shown in Fig. 10 approaches a maximum value asymptotically because this analysis includes the off-line cost for setting up the emulator. For applications of RBMs to heavier-mass nuclei (discussed in Sec. V.B), which are also significantly more costly to solve for computationally using high-fidelity models, speed-up factors of the order of 10^6–10^9 have been observed.

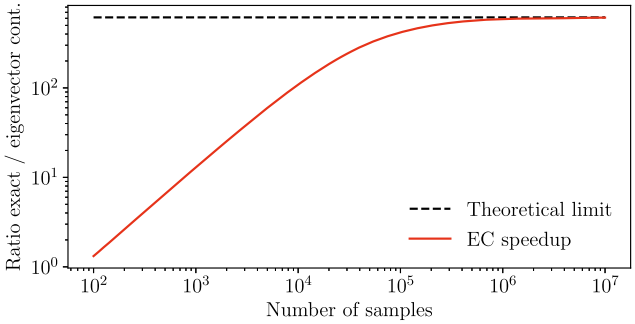

FIG. 10. Speedup factor (ratio of estimated required floating-point operations) of EC emulation compared to direct calculation as a function of the number of samples, i.e., the number of calls to the emulator. The theoretical limit indicates the maximum speedup reached asymptotically (as the off-line cost becomes amortized) in the number of samples, which is 614 in this example. Adapted from König *et al.*, 2020.

Wesolowski *et al.* (2021) and Djärv *et al.* (2022) used EC to construct fast and accurate emulators of no-core shell-model calculations in the analysis of three-nucleon forces in χEFT. Djärv *et al.* (2022) analyzed ^{6}Li in the m scheme, i.e., laboratory coordinates. Thus, the dimensionality of the Hamiltonian matrix also grew dramatically. For $N_{\max} = 8$ the matrix dimension in the large space is already $N_h \times N_h = 10^6 \times 10^6$, which requires significant computational efforts to be diagonalized even when using the Lanczos method. Becker *et al.* (2023) expressed the Hamiltonian in a symplectic symmetry-adapted no-core shell-model basis and used EC to further reduce the dimensionality of the Hamiltonian to construct accurate emulators for ^{12}C.

B. Subspace-projected coupled cluster

In nuclear physics and quantum chemistry one often encounters matrix representations of the many-body Schrödinger equation that are too large to permit an efficient diagonalization. For such cases, the coupled cluster (CC) method (Shavitt and Bartlett, 2009) can be an effective tool for approximating solutions in a space with dimensionality that is significantly reduced compared to an asymptotically exact method such as the no-core shell model. CC is based on a similarity-transformed Hamiltonian $\bar{H}(\boldsymbol{\theta}) = e^{-T} H(\boldsymbol{\theta}) e^{T}$, where the cluster operator $T = T_1 + T_2 + \cdots + T_n + \cdots + T_A$ is the sum of n-particle n-hole excitation (np-nh) operators acting on an A-particle vacuum state $|\Phi\rangle$. In nuclear physics, T is usually truncated at the coupled cluster singles and doubles (CCSD) level, i.e., $T = T_1 + T_2$, with triples excitations T_3 incorporated perturbatively (Hagen *et al.*, 2014). This typically captures 99% of the correlation energy of closed (sub)shell systems (Bartlett and Musiał, 2007; Ekström *et al.*, 2023a).

Truncating T makes the similarity transformation nonunitary and $\bar{H}(\boldsymbol{\theta})$ non-Hermitian. In the CCSD approximation, the biorthogonal left and right eigenstates of $\bar{H}$ are obtained as solutions to the CCSD equations, which can be viewed as a set of Galerkin equations based on $1p$-$1h$ and $2p$-$2h$ test functions $|\Phi_i^a\rangle \equiv a_a^\dagger a_i|\Phi\rangle$ and $|\Phi_{ij}^{ab}\rangle \equiv a_a^\dagger a_b^\dagger a_j a_i|\Phi\rangle$. CC calculations of atomic nuclei belong to a class of *ab initio* methods that scale polynomially with system size. Nevertheless, high-fidelity and state-of-the-art calculations beyond the lightest-mass nuclei require significant high-performance computing resources. Ekström and Hagen (2019) extended EC to non-Hermitian Hamiltonian matrices and the CC method. This has paved the way for fast and accurate emulations of properties of atomic nuclei and sophisticated computational statistics analyses.

Subspace-projected CC (SPCC) is an RBM using snapshots of the bivariational left and right CC states obtained at n_b different values of the parameters $\boldsymbol{\theta}$, such as the coupling constants in the description of the strong interaction Hamiltonian. The matrix elements of the $n_b \times n_b$ SPCC Hamiltonian and norm matrices have been worked out for the case of reference states built from harmonic-oscillator single-particle states (Ekström and Hagen, 2019). Typically CC calculations exploit a Hartree-Fock reference state. This embeds a dependence on $\boldsymbol{\theta}$ in the basis states that makes the evaluation of the matrix elements of the SPCC Hamiltonian and the norm matrices more cumbersome but can be done using a generalized Wick's theorem. The use of SPCC with more complex reference states is currently being explored. Inspired by the success of SPCC, an RBM of an angular-momentum-projected Hartree-Fock approximation was recently applied to emulate Hartree-Fock calculations of excited states in axially deformed nuclei (Ekström *et al.*, 2023b).

The first application (Ekström and Hagen, 2019) of SPCC to an atomic nucleus (^{16}O) demonstrated that $n_b \approx 50$ CCSD snapshots are sufficient to accurately emulate, i.e., with subpercent precision, realistic predictions of the energy and charge radius of the ground state in this nucleus as a function of $\boldsymbol{\theta}$. Here $\boldsymbol{\theta}$ denotes the 16 low-energy constants of a nuclear interaction description at next-to-next-to-leading order in χEFT. The accuracy of the SPCC emulator is noteworthy even when one uses few snapshots in a wide range of values for $\boldsymbol{\theta}$. Indeed, $n_b = 64$ snapshots $\boldsymbol{\theta}_i$ in a Latin hypercube design covering an extremely dispersed set of predictions for the energy and radius in ^{16}O is sufficient to obtain ∼97% accuracy compared to exact CCSD predictions; see Fig. 11. Narrowing the set of snapshots to a physically motivated parameter domain increases the accuracy significantly while using even fewer snapshots. SPCC emulators are also typically fast, and the bulk properties of ^{16}O could be sampled for 10^6 values of $\boldsymbol{\theta}$ in one hour on a standard laptop computer, while an equivalent set of exact CCSD calculations would require 20 years of single-node compute time, i.e., an observed speedup factor of 10^5.

Since the first application, SPCC has been extended to emulate the properties of ground and excited states in heavier-mass nuclei (Hu *et al.*, 2022; Kondo *et al.*, 2023) and infinite nuclear matter (Jiang *et al.*, 2024a) at different levels of fidelity up to perturbative triples excitations. The low computational cost of the SPCC method, with observed speedup factors of 10^9, has thus enabled a wide range of exciting computational statistics analyses expounding on how nuclear properties are linked to effective field theory descriptions of the strong interaction.

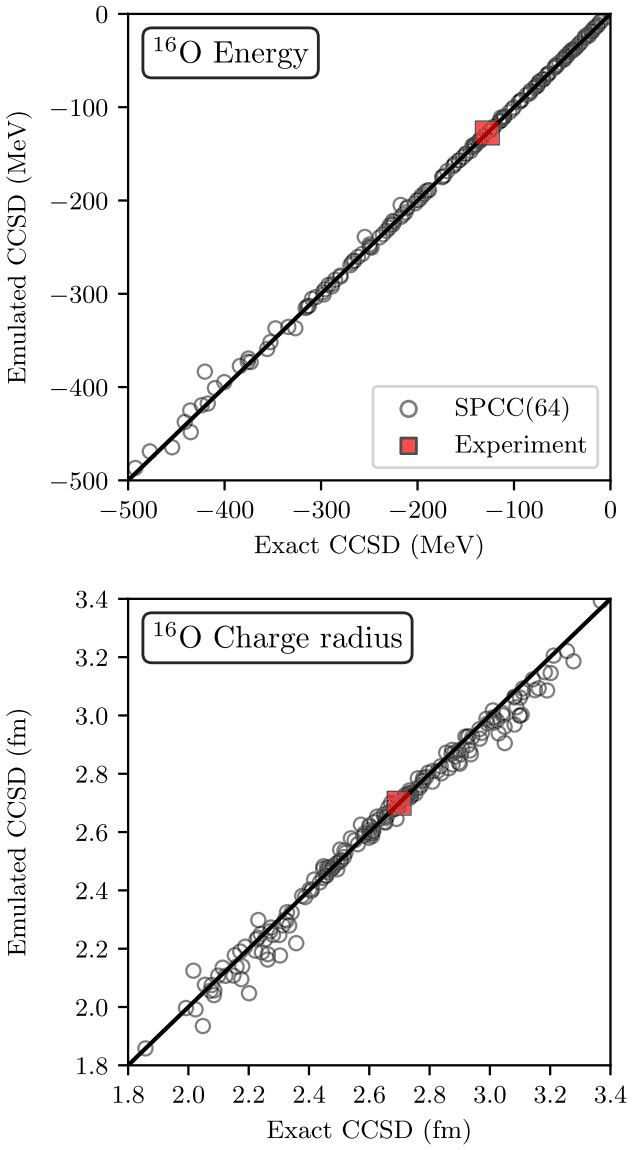

FIG. 11. Comparison of SPCC, based on $n_b = 64$ snapshots, and exact CCSD calculations for (top panel) the ground-state energy and (bottom panel) charge radius across a wide range of values. Adapted from Ekström and Hagen, 2019.

The CC method follows a bivariational principle that renders the SPCC Hamiltonian non-Hermitian. This may lead to difficulties identifying the target state in the spectrum as it is not guaranteed to be the lowest. Indeed, in heavier-mass nuclei and nuclear matter, where level densities are higher, the target state can sometimes appear to be an excited state. The interpretation of the SPCC spectrum under these conditions remains an open challenge. Using the bivariational principle, Jiang *et al.* (2024b) introduced a method called small-batch voting to detect the target state in such scenarios. Much remains to be discovered regarding the convergence properties

of non-Hermitian EC and the advantages of RBMs applied to the CC method.

C. Phenomenological shell model

While emulating low-energy constants stemming from χEFT is probably one of the most relevant scenarios in *ab initio* nuclear structure theory, EC provides opportunities to enhance phenomenological approaches as well. One such application was given by Yoshida and Shimizu (2022), who applied EC in connection with the nuclear shell model. In a shell-model calculation, the Hamiltonian is typically split into one- and two-body terms $H = H^{(1)} + H^{(2)}$, with natural extensions to include higher-body terms. $H^{(1)}$ models the effective mean field that generates nuclear orbitals, while $H^{(2)}$ describes interaction among valence nucleons. Input parameters for a shell-model calculation are single-particle energies that determine the diagonal part of $H^{(1)}$ and two-body matrix elements that parametrize the interaction in $H^{(2)}$. Both types of parameters need to be fitted to experimental data within the regime of nuclei that one wants to describe with a particular model (for example, *sd* shell nuclei).

As a typical starting point, Yoshida and Shimizu (2022) considered the unified *sd* shell interaction B (USDB) for *sd* shell nuclei (Brown and Richter, 2006), which spans overall a 66-dimensional parameter space of three single-particle energies and 63 two-body matrix elements, collected into a vector $\boldsymbol{\theta}$. While each individual diagonalization within the valence space can be economical, the large number of parameters implies that there is significant potential for speeding up the fitting process via EC emulation. Setting up such an emulator follows the standard EC procedure based on training points $\boldsymbol{\theta}_i$ with $i = 1, \ldots, n_b$ together with the Hamiltonian

$$\begin{aligned}\tilde{H}_{ij} &= \langle\psi(\boldsymbol{\theta}_i)|H(\boldsymbol{\theta}_*)|\psi(\boldsymbol{\theta}_j)\rangle \\ &= \sum_k h_k^{(1)} \times \overline{\mathrm{OBTD}}_k + \sum_k V_k^{(2)} \times \overline{\mathrm{TBTD}}_k, \end{aligned} \tag{21}$$

where $h_k^{(1)}$ and $V_k^{(2)}$ are the single-particle energies and two-body matrix elements that multiply one- and two-body transition densities ($\overline{\mathrm{OBTD}}$ and $\overline{\mathrm{TBTD}}$), respectively, and $|\psi(\boldsymbol{\theta}_i)\rangle$ denotes a particular shell-model wave function obtained for parameters $\boldsymbol{\theta}_i$. Equation (21) has an affine structure that enables the previously mentioned on-line–off-line decomposition.

A number of benchmark scenarios are considered, varying the number of training points n_b from 50 to 250, and also the number of states (lowest part of the spectrum starting with the ground state) per training parameter set, from 1 to 5. Overall, for a selection of nuclei such as ^{28}Si and ^{24}Mg relative emulator errors of the order between less than 1% up to a few percent are observed. A Monte Carlo sampling technique is proposed to assign emulator uncertainties for individual evaluations.

In line with other work such as König *et al.* (2020), Yoshida and Shimizu (2022) found that emulated wave functions generally show larger emulation discrepancies than binding energies, leading to a larger spread for emulator evaluations of operators such as magnetic dipole moments and quadrupole moments. To improve the emulator accuracy and avoid problems in correctly describing such observables, Yoshida and Shimizu (2022) suggested using the shell-model emulator as a preprocessor to generate optimized initial states for a subsequent exact Lanczos diagonalization. More generally there will be challenges in applying EC beyond the *sd* shell (for example, to the *pf* shell), where there are many more parameters and the time to generate low-lying states in the off-line training phase will be greater. This is where the experience from the RBM community in reduced-order sampling (such as using SVD methods; see Sec. III) could be profitably carried over to nuclear problems.

VI. EXAMPLES OF EXTENSIONS

In Secs. VI.A–VI.C we introduce three extensions of the basic EC method.

A. Emulators for quantum scattering

Uncertainty quantification will often require calculations of scattering observables with many different Hamiltonian parametrizations. Examples in nuclear physics include the calibration of χEFT interactions and of phenomenological optical potentials. This has led to the extension of model-driven emulators to the quantum mechanical two-body scattering problem and beyond. Of particular importance for nuclear applications is the ongoing development of three-body scattering emulators.

Quantum scattering is not an eigenvalue problem, but the same principles that make EC effective for bound states carry over to scattering. In the time-independent formulation of scattering, we still start with the strong form of the Schrödinger equation $H(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle = E|\psi(\boldsymbol{\theta})\rangle$, but now E is specified rather than determined (although it can also be treated as a parameter of the emulator). The freedom to formulate the Schrödinger equation for scattering in different ways, including homogeneous or inhomogeneous differential equations for scattering wave functions as well as integral equations for scattering matrices, leads to many possible emulators. In addition, there is the freedom to choose trial and test bases for Galerkin projection; see Sec. III. As such, beyond the intrinsic use of scattering emulators the scattering problem is a prototype for multiple approaches to model reduction in other settings.

There are numerous variational formulations of scattering, such as those due to Kohn, Schwinger, and Newton (Newton, 2002). Variational here means that there is a stationary functional, but in most cases this does not imply that the result is an upper bound, unlike the case with bound states. For each of these variational formulations there is a corresponding RBM emulator.

The first implementation of a quantum scattering emulator (Furnstahl *et al.*, 2020) used the Kohn variational principle (KVP) for partial wave scattering (Kohn, 1948). For two-body scattering in a single channel with angular momentum l at on-shell energy $E = q^2/2\mu$, the KVP functional takes the form

$$\mathcal{K}[\tilde{\psi}] = \tilde{K}_E + \langle\tilde{\psi}|H - E|\tilde{\psi}\rangle. \tag{22}$$

In Eq. (22) the trial scattering wave function $|\tilde{\psi}\rangle$ in position space is constrained to satisfy the asymptotic normalization condition

$$\tilde{\psi}_l(r) \underset{r\to\infty}{\longrightarrow} j_l(qr) + n_l(qr)\tan\delta_l, \tag{23}$$

and

$$\tilde{K}_E = -\frac{\tan\delta_l}{2\mu q} \tag{24}$$

is the on-shell K matrix corresponding to the phase shift $\delta_l(E)$. This functional is stationary about the exact solution ψ such that $\mathcal{K}[\psi+\delta\psi] = K_E + \mathcal{O}(\delta K)^2$.

An EC-RBM emulator for the KVP uses a snapshot trial basis as in Eq. (10), where each basis wave function satisfies Eq. (23) and the overall constraint for the trial wave function, which requires $\sum_{i=1}^{n_b}\beta_i = 1$, is enforced by a Lagrange multiplier. Varying the KVP functional with this constraint yields a low-dimensional $(n_b \times n_b)$ linear matrix problem. If the Hamiltonian is affine in the parameters, all of the relevant matrix elements can be precomputed in the off-line stage as in Fig. 4. An example of this emulator from Furnstahl *et al.* (2020) is shown in Fig. 12 for a model nucleon-nucleon potential with two parameters (the strengths of two Gaussians). The snapshot wave functions for four randomly chosen sets of $\boldsymbol{\theta}_i$ are shown in Fig. 12(a), while the corresponding phase shifts are shown in Fig. 12(b). Despite no indication from Fig. 12 that this is a good basis, the emulator is fast and accurate through the full range of energies.

The KVP is sometimes itself used as a high-fidelity solution method, where it is well known to be plagued with numerical issues known as Kohn anomalies. These can be mitigated for emulators by a more general formulation than Eqs. (22)–(24) that uses multiple scattering matrices (rather than just the K matrix); see Drischler *et al.* (2021, 2023) for details. This approach was extended to coupled channels and to momentum space by Garcia *et al.* (2023), with successful tests of the full range of two-body scattering observables using a state-of-the-art χEFT Hamiltonian with 25 parameters (up to six in each partial wave channel, with the parameters emulated independently). Speedups of 2 orders of magnitude over high-fidelity calculations were found even when using basis sizes large enough to achieve a mean relative emulator error of the order of 10^{-10} over a wide region in parameter space (in practice this means that n_b is equal to twice the number of parameters in a given channel).

Another form of the KVP-type emulator avoids using a Lagrange multiplier to constrain the normalization of basis wave functions by introducing a trial basis only for the second (scattering) term of Eq. (23) rather than for the full wave function. The free wave function [the first term in Eq. (23)] fixes the normalization. The Schwinger and Newton emulators use alternative variational principles, with the latter having a trial basis of K matrices rather than wave function (Melendez *et al.*, 2021); it is applied to the calibration of χEFT parameters given by Svensson, Ekström, and Forssén (2024).

Each of these variational formulations has a Galerkin counterpart, so we can use a Galerkin projection as an alternative path to constructing the emulators. This is worked out for each of the Kohn, Schwinger, and Newton emulators given by Drischler *et al.* (2023). This also means that we can directly formulate scattering emulators that do not have an obvious variational counterpart. With the normalization fixed at the origin ($r = 0$) by a free solution value and first derivative, the snapshot basis of scattering terms can be used in a Galerkin projection of Eq. (12). Application of this emulator to calibrate phenomenological optical potentials by Odell *et al.* (2024), implemented with ROSE software from the BAND project (Beyer *et al.*, 2023), uses proper orthogonal decomposition (see Sec. III.A) to optimize the basis and demonstrates a method to handle the nonaffine parameters of the potential. Yet another formulation builds on R-matrix theory with successful applications to fusion observables (Bai and Ren, 2021; Bai, 2022). The frontier for scattering emulators is for three-body problems. A proof-of-principle demonstration using the KVP for three bosons was given by Zhang and Furnstahl (2022), and tests of realistic nuclear scattering are in progress.

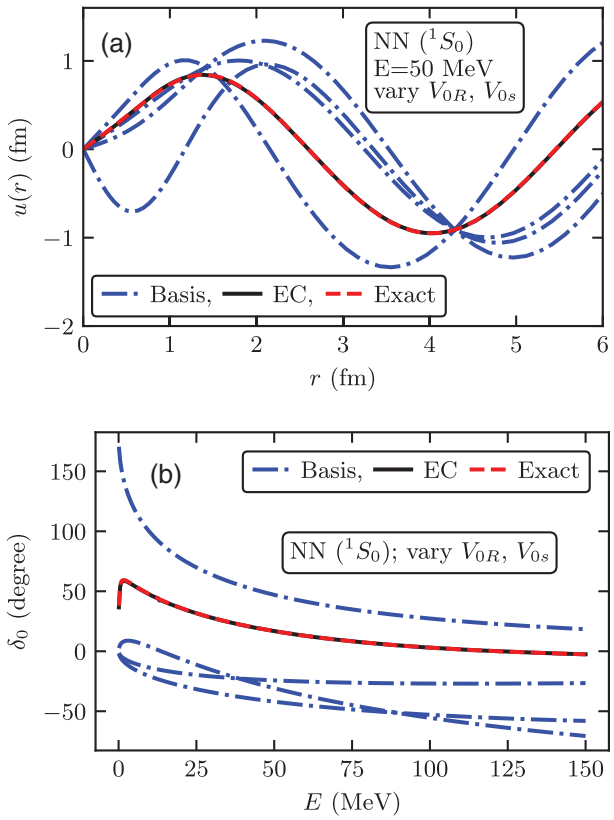

FIG. 12. (a) Scattering wave functions for a model nucleon-nucleon potential at a fixed energy. The dot-dashed curves are for four choices of $\boldsymbol{\theta}_i = \{V_{0R}, V_{0s}\}$ that compose the trial basis, the dashed curve is for the target values, and the solid curve is the prediction using the KVP emulator. The curves have a common crossing point at the value of r where the second term in Eq. (23) is zero. (b) Scattering phase shifts for the same parameter sets and the emulator prediction. From Furnstahl *et al.*, 2020.

B. Finite-volume dependence and resonances

Another extension of EC, developed by Yapa and König (2022), is concerned with extrapolating or interpolating the

volume dependence of energy levels in finite periodic boxes, with the particular application of studying resonance properties via finite-volume (FV) simulations (Wiese, 1989; Lüscher, 1991; Rummukainen and Gottlieb, 1995; Klos *et al.*, 2018). In this scenario, not only does the Hamiltonian $H = H(L)$ depend explicitly on the size of a cubic box L (via the periodic extension of the interaction part) but, since eigenstates of $H(L)$ have to satisfy the periodic boundary condition, they also carry an implicit dependence on L. Specifically, states defined in boxes with different L are vectors in distinct Hilbert spaces, which makes it *a priori* difficult to give a well-defined meaning to matrices

$$\tilde{H}_{ij}(L_*) = \langle \psi_{L_i} | H(L_*) | \psi_{L_j} \rangle, \tag{25a}$$

$$\tilde{N}_{ij} = \langle \psi_{L_i} | \psi_{L_j} \rangle \tag{25b}$$

that appear in a standard EC setup, i.e., a Galerkin projection with snapshots defined with training points L_i and a target volume L_*.

Yapa and König (2022) overcame this issue by defining a space

$$\mathcal{H} = \bigcup_{\{L>0\}} \mathcal{H}_L. \tag{26}$$

as a union of Hilbert spaces $\mathcal{H}_L$ that contain periodic states with a fixed period L. This set $\mathcal{H}$ is not a vector space with the standard pointwise addition of functions (assuming that $\mathcal{H}_L$ are simple function spaces), but it can be made into one by defining appropriate operations that combine functions with different periods.

Yapa and König (2022) accomplished this by applying dilatations (transformations involving stretching and rescaling) that map states onto a common space prior to applying the standard operations within that space and showed that, when this procedure was applied to a truncated bases of periodic functions [such as simple plane waves or the discrete variable representation used by Klos *et al.* (2018) to study few-body systems in finite volume], ultimately the outcome is equivalent to simply operating within the $\mathbb{R}^n$ space of coefficient vectors (where n denotes the dimension of the finite space). We note that, in the broader RBM context, problems such as the volume dependence discussed here have been treated by mapping the physical domain to a fixed reference domain and formulating an equivalent problem on this reference domain (Rozza, 2005; Rozza and Veroy, 2007).

As mentioned at the outset of this section, studying resonances in FV was a primary motivation for the development of finite-volume eigenvector continuation (FVEC). Figure 13, from Yapa and König (2022), demonstrates this application with the example of a three-boson resonance generated by a sum of attractive and repulsive Gaussian potentials (Blandon, Kokoouline, and Masnou-Seeuws, 2007). For almost the entire range of volumes shown in Fig. 13, FVEC produces results that in the resolution of the plot are virtually indistinguishable from exact calculations, and the avoided level crossing around $L \sim 28$ fm, indicating the resonance, is well reproduced.

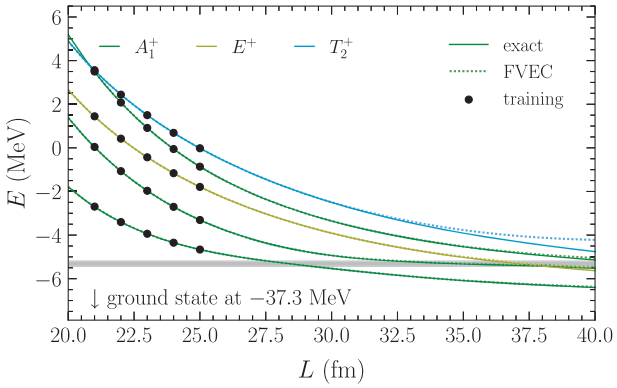

FIG. 13. Positive-parity finite-volume energy spectrum of three bosons exhibiting a resonance state. The solid lines show the exact states calculated in a certain basis of the discrete variable representation, whereas the dashed lines indicate FVEC results obtained based on training data at five different box sizes (solid circles) using the eight lowest states in the spectrum at each volume (including the ground state, which is not shown). From Yapa and König, 2022.

Resonances are a fascinating phenomenon found in many areas of physics that are closely related to the study of open quantum systems. Numerically studying their properties is notoriously challenging because accommodating states that decay with a finite lifetime requires either a time-dependent treatment or special "tricks" to describe them within a time-independent framework. Enclosing the system in a finite volume and looking for avoided level crossing in the volume-dependent energy spectrum is an elegant way to identify resonances, but this approach is geared primarily toward few-body systems. In formal scattering theory, decaying resonances are generally associated with poles of the scattering matrix (the "S matrix") at complex energies $E = E_R - i\Gamma/2$, located in the fourth quadrant of the complex plane. The real part E_R denotes the resonance position, while the width $\Gamma > 0$ is related to the inverse of the lifetime.

While ordinary Hermitian quantum mechanics can only describe either bound states (real $E < 0$) or scattering (real $E > 0$), different options to achieve non-Hermitian extensions have been developed in order to allow for complex energy eigenvalues. Yapa, Fossez, and König (2023) developed an extension of EC that uses the so-called uniform complex scaling technique to describe resonances, and, in particular, their trajectories in the complex plane under variation of the Hamiltonian, which is written as $H = H(\theta)$. Complex scaling is based on rotating radial coordinates according to $r \to re^{i\phi}$ with an angle $\phi > 0$ or, equivalently (Afnan, 1991), the conjugate momentum variable q according to $q \to qe^{-i\phi}$. Along this rotated contour, resonance wave functions behave effectively like bound states, and the complex-scaled (rotated) Hamiltonian allows for complex energy eigenvalues. An important aspect of complex scaling (as well as other methods that enable the description of resonance in time-independent quantum mechanics) is that inner products of complex-scaled states do not involve complex conjugation of the "bra" state. Yapa, Fossez, and König (2023) showed that eigenvector continuation for resonance states can be implemented by

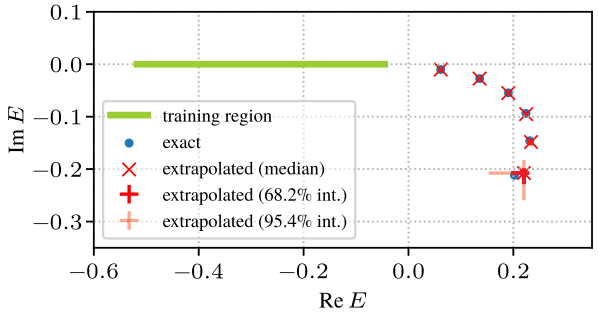

FIG. 14. Bound-state-to-resonance extrapolation performed with CAEC for a two-body system supporting an S-wave resonance for certain values of the parameter c. Five training points were randomly drawn from the region $c \in (0.9, 1.3)$ per dataset, leading to bound states within the shaded line along the negative real axis. Multiple samples of these five points were used to obtain extrapolations with uncertainty estimates. From Yapa, Fossez, and König, 2023.

defining the Hamiltonian and norm matrix elements in terms of the so-called c product (Moiseyev, Certain, and Weinhold, 1978; Moiseyev, 2011), which for eigenstates $|\psi_1\rangle$ and $|\psi_2\rangle$ with equal angular-momentum quantum numbers is given by

$$\langle \psi_1 | \psi_2 \rangle = \int dr\, \psi_1(r)\psi_2(r). \tag{27}$$

Standard EC works well in this way for extrapolating (or interpolating) the trajectory of a resonance state as it moves in the fourth quadrant as a function of θ. While this is relevant, for example, for constructing EC-based emulators for resonance properties, Yapa, Fossez, and König (2023) are furthermore interested in the case where EC is trained within a regime of θ, where the state is actually bound, and then extrapolate from there into the resonance domain. The key result of that work is that while using the c product alone is not sufficient to achieve this, an extension of EC that also includes for each training bound state its complex conjugate (with complex scaling, bound-state eigenvalues remain real, but the corresponding wave functions defined along the rotated contour have nontrivial complex behavior). This conjugate-augmented eigenvector continuation (CAEC) is then able to perform the desired extrapolation from bound states to resonances. An example of this is shown in Fig. 14. While Yapa, Fossez, and König (2023) considered as proof of concept only two-body resonances calculated with complex scaling, they conjectured that generally CAEC is expected to work for quantum systems involving more particles, as well as in conjunction with techniques other than complex scaling.

C. Quantum Monte Carlo simulations

Quantum Monte Carlo simulations are widely used for first-principles calculations of quantum many-body systems across many subfields of physics. In cases where sign oscillations are not a problem, the computational effort usually scales as a low-order polynomial in the number of particles. Since quantum Monte Carlo simulations can work with vectors in extremely large linear spaces, the combination of EC with quantum Monte Carlo methods is potentially powerful. The application of EC with quantum Monte Carlo was discussed by Frame *et al.* (2018) as well as by Frame (2019).

In EC we perform a Galerkin projection and thus need to compute inner products between energy eigenstates associated with different Hamiltonians. However, computing the inner product of the different eigenstates is not straightforward using quantum Monte Carlo simulations. We illustrate the problem with an example involving ground-state wave functions. Let H_A and H_B be two quantum Hamiltonians with ground-state wave functions $|v_A^0\rangle$ and $|v_B^0\rangle$, respectively, and ground-state energies E_A^0 and E_B^0, respectively. Let $|\phi\rangle$ be any state that is not orthogonal to $|v_A^0\rangle$ and $|v_B^0\rangle$. Starting with the state $|\phi\rangle$, we can obtain $|v_A^0\rangle$ by applying the Euclidean time evolution operator $e^{-H_A t}$ and taking t to be large and positive. Similarly, we can obtain $|v_B^0\rangle$ by applying the Euclidean time evolution operator $e^{-H_B t}$. In the limit of large t, we have

$$e^{-H_A t}|\phi\rangle \approx e^{-E_A^0 t}\langle v_A^0|\phi\rangle |v_A^0\rangle, \tag{28}$$

$$e^{-H_B t}|\phi\rangle \approx e^{-E_B^0 t}\langle v_B^0|\phi\rangle |v_B^0\rangle. \tag{29}$$

The difficulty arises from the fact that $|v_A^0\rangle$ and $|v_B^0\rangle$ appear with exponential factors of $e^{-E_A^0 t}$ and $e^{-E_B^0 t}$, respectively. Calculations of the magnitude of the inner product $\langle v_A^0|v_B^0\rangle$ are prone to large relative errors since the amplitude is dominated by factors of $e^{-E_A^0 t}$ and $e^{-E_B^0 t}$ for large t.

A technique called the floating block method that addresses this problem was introduced by Sarkar, Lee, and Meißner (2023). The floating block method is based on the observation that

$$\lim_{t\to\infty} \frac{\langle\phi|e^{-H_A t}e^{-H_B t}e^{-H_A t}e^{-H_B t}|\phi\rangle}{\langle\phi|e^{-2H_A t}e^{-2H_B t}|\phi\rangle} = |\langle v_A^0|v_B^0\rangle|^2. \tag{30}$$

We note that the problematic exponential factors of $e^{-E_A^0 t}$ and $e^{-E_B^0 t}$ cancel from this ratio. We can also calculate the complex phase of the inner product using

$$\lim_{t\to\infty} \frac{\langle\phi|e^{-2H_A t}e^{-2H_B t}|\phi\rangle}{|\langle\phi|e^{-2H_A t}e^{-2H_B t}|\phi\rangle|} = \frac{\langle v_A^0|v_B^0\rangle}{|\langle v_A^0|v_B^0\rangle|}. \tag{31}$$

In Eq. (31) we are using the phase convention that $\langle v_A^0|\phi\rangle$ and $\langle v_B^0|\phi\rangle$ are positive.

If we try to compute the ratio of the numerator and denominator in Eq. (30) directly using Monte Carlo simulations, the result will still be noisy since the numerator and denominator are uncorrelated with each other. To overcome this problem, the floating block method instead computes ratios of quantities that are strongly correlated. We define $Z(t_1, t_2, t_3, t_4)$ to be the amplitude

$$Z(t_1, t_2, t_3, t_4) = \langle\phi|e^{-H_A t_1}e^{-H_B t_2}e^{-H_A t_3}e^{-H_B t_4}|\phi\rangle. \tag{32}$$

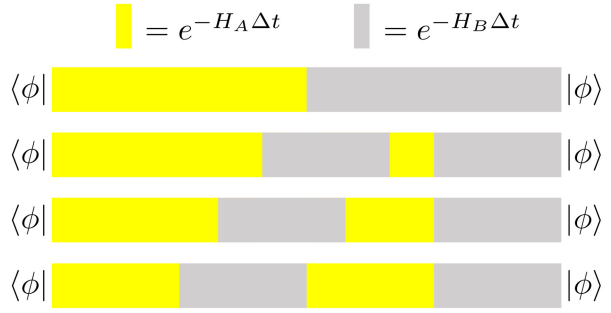

FIG. 15. Schematic showing the intermediate values for the Euclidean time blocks, gradually interpolating between $\langle\phi|e^{-2H_A t}e^{-2H_B t}|\phi\rangle$ in the denominator of Eq. (30) and $\langle\phi|e^{-H_A t}e^{-H_B t}e^{-H_A t}e^{-H_B t}|\phi\rangle$ in the numerator of Eq. (30) (Sarkar, Lee, and Meißner, 2023). This corresponds to gradually changing the values of t_1, t_2, t_3, and t_4 in Eq. (32).

In the floating block method, we compute ratios of the form

$$\frac{Z(t_1,t_2,t_3,t_4)}{Z(t_1',t_2',t_3',t_4')} \tag{33}$$

for values t_1, t_2, t_3, t_4 and t_1', t_2', t_3', t_4' that are close to each other. We can then form telescoping products of such ratios,

$$\frac{Z(t_1,t_2,t_3,t_4)}{Z(t_1',t_2',t_3',t_4')}\frac{Z(t_1',t_2',t_3',t_4')}{Z(t_1'',t_2'',t_3'',t_4'')}\frac{Z(t_1'',t_2'',t_3'',t_4'')}{Z(t_1''',t_2''',t_3''',t_4''')}\cdots. \tag{34}$$

In this manner, we can calculate the ratio of the numerator and denominator in Eq. (30). This is illustrated in Fig. 15.

Sarkar, Lee, and Meißner (2023) used the floating block method to compute the binding energies of ^{4}He, ^{8}Be, ^{12}C, and ^{16}O using Monte Carlo simulations with a lattice Hamiltonian of the form $H_{\text{free}} + c_L V_L + c_{\text{NL}} V_{\text{NL}}$. V_L is a two-nucleon interaction with local interactions, meaning that the interaction does not move the relative positions of the nucleons. V_{NL} is a two-nucleon interaction composed of nonlocal interactions where the relative positions of the nucleons are allowed to change. Sarkar, Lee, and Meißner (2023) normalized V_L and V_{NL} so that $(c_L, c_{\text{NL}}) = (1,0)$ and $(c_L, c_{\text{NL}}) = (0,1)$ both gave realistic results for ^{4}He. In Fig. 16, we plot the ground-state energy of ^{16}O relative to the four-alpha threshold, $E(^{16}\text{O}) - 4E(^4\text{He})$ (Sarkar, Lee, and Meißner, 2023). The EC calculation is performed with $n_b = 2$ snapshots at $(c_L, c_{\text{NL}}) = (0.5, 0.5)$ and $(0,1)$ in a periodic box of length $L = 15.76$ fm. The dashed line shows the contour for the observed experimental value. The zero contour line corresponds with the quantum phase transition where ^{16}O falls apart into four alpha particles. These results are consistent with the finding of Elhatisari *et al.* (2016) that, without sufficiently attractive local interactions, symmetric nuclear matter forms a Bose gas of alpha particles rather than a nuclear liquid.

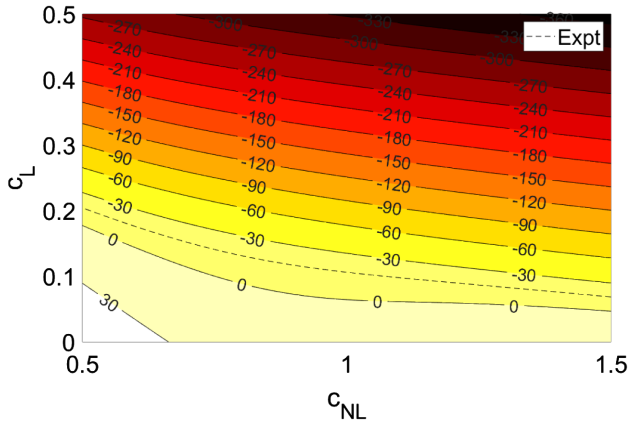

FIG. 16. Contour plots for the difference between the EC emulated energy for ^{16}O and the four-alpha threshold energy $E(^{16}\text{O}) - 4E(^4\text{He})$. c_L is the coefficient of the local two-nucleon interaction, and c_{NL} is the coefficient of the nonlocal two-nucleon interaction. The dashed line shows the contour for the observed experimental value.

VII. SUMMARY AND FUTURE DIRECTIONS

In this Colloquium we have presented the historical development, the theoretical framework, and the applications of EC and projection-based emulators. The key concept is that the eigenvector $|\psi(\boldsymbol{\theta})\rangle$ is an analytic function for real values of the parameters and approximately lies on a linear subspace with a finite number of dimensions. The smoother and more gradual the undulations, the fewer dimensions needed. The linear subspace can be found efficiently by taking snapshots of $|\psi(\boldsymbol{\theta}_i)\rangle$ for the selected parameter values $\boldsymbol{\theta}_i$ and using the corresponding subspace spanned by the snapshots.

EC is part of a larger class of subspace-projection techniques called RBMs, which are themselves part of a yet larger category of model-driven reduced-order models. The development of EC has emphasized applications to quantum systems, from few-body problems to many-body problems, and from bound states to scattering states and resonances. Some of the addressed topics go beyond the traditional class of problems typically encountered in the reduced-basis literature, such as parameter extrapolation to domains that are not directly calculable, accelerating the convergence of many-body perturbation theory, and working with extremely large or infinite-dimensional vector spaces. As noted in Sec. IV, EC can sometimes face challenges when it is applied to situations where one tries to extrapolate across boundaries between physically distinct phases.

While the development of EC and projection-based emulators by the nuclear theory community has naturally focused on problems of interest in nuclear physics, such as uncertainty quantification and other computational statistics analyses of the nuclear Hamiltonian, the methods are general and can be applied to other fields where quantum wave functions are important. Remaining challenges include how to identify EC target states in the spectrum of non-Hermitian Hamiltonians and how to best handle nonaffine parameter dependencies in nuclear applications. The usefulness of combining distributed emulators as miniapplications, for example, in Bayesian inference analyses, also has yet to be capitalized on.

Other areas where EC and projection-based emulators should be useful are atomic and molecular physics, ultracold atomic gases, strongly correlated electronic systems, quantum

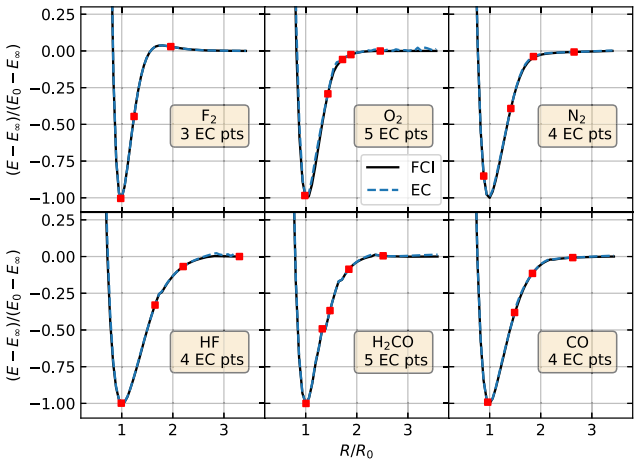

FIG. 17. The potential energy vs the bond stretching factor R/R_0 for F_2, O_2, N_2, HF, H_2CO, and CO obtained using EC with several snapshots. We show the comparison with exact FCI calculations. From Mejuto-Zaera and Kemper, 2023.

spin liquids, and quantum chemistry. Mejuto-Zaera and Kemper (2023) applied EC to problems in *ab initio* quantum chemistry. In Fig. 17, we show the potential energy surface for several molecules versus the bond stretching factor R/R_0. The molecules are F_2, O_2, N_2, HF, H_2CO, and CO. We show the EC results obtained with $n_b = 3$, 4, or 5 snapshots and the comparison with the exact full configuration interaction (FCI) calculations. Eigenvector continuation is working well in reproducing all of the potential energy surfaces.

The potential energy surface calculations described by Mejuto-Zaera and Kemper (2023) can also be performed on a quantum computer, and the corresponding algorithm is called quantum EC (Francis *et al.*, 2022). While much effort in the quantum computing community has focused on variational methods optimizing a single trial vector, variational calculations using subspace projection can in principle provide a better approximation to the eigenstate of interest for the same computational resources.

In quantum EC, the same general approach is used as with a classical computer, though there are some technical differences. On a digital quantum computer with N qubits, we start with the state where all qubits are in the $|0\rangle$ state: $|00\cdots0\rangle$. Let $|\psi_i\rangle$ denote eigenstate snapshots at parameter values $\boldsymbol{\theta}_i$. We assume that our chosen quantum eigenstate algorithm gives us some unitary gate U_i such that the action on $|00\cdots0\rangle$ gives us a good approximation to $|\psi_i\rangle$.

The norm and Hamiltonian matrix elements can be determined using an ancilla qubit. We use the ancilla qubit to apply the controlled operations for U_i and $U_j^\dagger$. Such a controlled operation means that we perform the transformation only if the ancilla qubit is in the $|1\rangle$ state. After these controlled unitary operators, we then apply σ_X or σ_Y rotations to the ancilla qubit and measure it, with either $|0\rangle$ or $|1\rangle$ as the outcome. This information is enough to determine the real and imaginary parts of $\langle\psi_j|\psi_i\rangle$. To compute the elements $\langle\psi_j|H(\boldsymbol{\theta})|\psi_i\rangle$, we decompose $H(\boldsymbol{\theta})$ into a sum of tensor products of Pauli operators. For each tensor product of Pauli operators U_P, we use the ancilla qubit to apply the controlled operations for U_i, U_P, and $U_j^\dagger$. We then apply σ_X and σ_Y rotations to the ancilla qubit and measure it (Francis *et al.*, 2022).

EC and projection-based emulators can also be combined with data-driven reduced-order model techniques such as Gaussian processes, neural networks, and dynamic mode decomposition. The combination of reduced-basis methods and machine learning is an active area of research and can be realized in many different ways. As noted in Sec. III, active learning methods (such as greedy algorithms) are often used to optimize the off-line selection of snapshot parameters and projection subspaces (Quarteroni, Manzoni, and Negri, 2016; Chellappa, Feng, and Benner, 2021; Sarkar and Lee, 2022). In numerically challenging problems where convergence with respect to the projection subspace dimension is slow, it is useful to treat the projection-based emulator as one component embedded within a larger framework such as a deep neural network or an autoencoder (Brunton and Kutz, 2019; Dal Santo, Deparis, and Pegolotti, 2020; Fresca *et al.*, 2020; Chen *et al.*, 2021).

ACKNOWLEDGMENTS

We are grateful to Christian Drischler, Pablo Giuliani, Gaute Hagen, and Xilin Zhang for reading the manuscript and suggesting improvements. The work of A. E. is supported by the European Research Council (ERC) under the European Union Horizon 2020 research and innovation program (Grant Agreement No. 758027) and the Swedish Research Council (Grant Agreement No. 2020-05127). The work of R. J. F. is supported by the U.S. National Science Foundation (NSF) [Grants No. PHY-2209442 and No. OAC-2004601 (CSSI BAND Collaboration)] and the U.S. Department of Energy (DOE) [Grants No. DE-SC0024509 (STREAMLINE Collaboration) and No. DE-SC0023175 (SciDAC-5 NUCLEI Collaboration)]. The work of S. K. is supported by the U.S. NSF (Grant No. PHY-2044632) and by the U.S. DOE [Grants No. DE-SC0024520 (STREAMLINE Collaboration) and No. DE-SC0024622]. This material is based upon work supported by the U.S. DOE, Office of Science, Office of Nuclear Physics, under the FRIB Theory Alliance, Award No. DE-SC0013617 (S. K.). The work of D. L. is supported by the U.S. NSF (Grant No. PHY-2310620) and the U.S. DOE [Grants No. DE-SC0013365, No.:DE-SC0023658, No. DE-SC0024586 (STREAMLINE Collaboration), and No. DE-SC0023175 (SciDAC-5 NUCLEI Collaboration)], and with computational resources provided by the Oak Ridge Leadership Computing Facility through the INCITE award "*Ab-initio* nuclear structure and nuclear reactions" and the GCS Supercomputer JUWELS at the Jülich Supercomputing Centre (JSC).

REFERENCES

Afnan, I. R., 1991, "Resonances in few body systems," Aust. J. Phys. **44**, 201–216.

Aktas, E., and F. Moses, 1998, "Reduced basis eigenvalue solutions for damaged structures," Mech. Struct. Mach. **26**, 63–79.

Bai, D., 2022, "New extensions of eigenvector continuation *R*-matrix theory based on analyticity in momentum and angular momentum," Phys. Rev. C **106**, 024611.

Bai, D., and Z. Ren, 2021, "Generalizing the calculable *R*-matrix theory and eigenvector continuation to the incoming wave boundary condition," Phys. Rev. C **103**, 014612.

Baran, V. V., and D. R. Nichita, 2023, "Reduced basis emulation of pairing in finite systems," Phys. Rev. B **107**, 144503.

Bartlett, R. J., and M. Musiał, 2007, "Coupled-cluster theory in quantum chemistry," Rev. Mod. Phys. **79**, 291–352.

Becker, K. S., K. D. Launey, A. Ekström, and T. Dytrych, 2023, "*Ab initio* symmetry-adapted emulator for studying emergent collectivity and clustering in nuclei," Front. Phys. **11**, 1064601.

Benner, P., M. Ohlberger, A. Patera, G. Rozza, and K. Urban, 2017, Eds., *Model Reduction of Parametrized Systems* (Springer, New York).

Benner, P., W. Schilders, S. Grivet-Talocia, A. Quarteroni, G. Rozza, and L. Miguel Silveira, 2020a, *Model Order Reduction, Vol. 2: Snapshot-Based Methods and Algorithms* (De Gruyter, Berlin).

Benner, P., W. Schilders, S. Grivet-Talocia, A. Quarteroni, G. Rozza, and L. Miguel Silveira, 2020b, *Model Order Reduction, Vol. 3: Applications* (De Gruyter, Berlin).

Benner, P., W. Schilders, S. Grivet-Talocia, A. Quarteroni, G. Rozza, and L. Miguel Silveira, 2021, *System- and Data-Driven Methods and Algorithms* (De Gruyter, Berlin).

Beyer, K., *et al.*, 2023, "BAND Framework: An open-source framework for Bayesian analysis of nuclear dynamics," Technical Report Version 0.3.0, https://github.com/bandframework/bandframework/releases/tag/v0.3.0.

Blandon, J., V. Kokoouline, and F. Masnou-Seeuws, 2007, "Calculation of three-body resonances using slow-variable discretization coupled with a complex absorbing potential," Phys. Rev. A **75**, 042508.

Bonilla, E., P. Giuliani, K. Godbey, and D. Lee, 2022, "Training and projecting: A reduced basis method emulator for many-body physics," Phys. Rev. C **106**, 054322.

Brehmer, P., M. F. Herbst, S. Wessel, M. Rizzi, and B. Stamm, 2023, "Reduced basis surrogates for quantum spin systems based on tensor networks," Phys. Rev. E **108**, 025306.

Brenner, S. C., and L. R. Scott, 2008, *The Mathematical Theory of Finite Element Methods*, Texts in Applied Mathematics Vol. 15 (Springer, New York).

Brink, D. M., and A. Weiguny, 1968, "The generator coordinate theory of collective motion," Nucl. Phys. **A120**, 59–93.

Broeckhove, J., and E. Deumens, 1979, "A mathematical foundation for discretisation techniques in the generator coordinate method," Z. Phys. A **292**, 243–247.

Brown, B. A., and W. A. Richter, 2006, "New 'USD' Hamiltonians for the *sd* shell," Phys. Rev. C **74**, 034315.

Brunton, S. L., and J. N. Kutz, 2019, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (Cambridge University Press, Cambridge, England).

Chellappa, S., L. Feng, and P. Benner, 2021, "A training set subsampling strategy for the reduced basis method," J. Sci. Comput. **89**, 1–34.

Chen, W., Q. Wang, J. S. Hesthaven, and C. Zhang, 2021, "Physics-informed machine learning for reduced-order modeling of nonlinear problems," J. Comput. Phys. **446**, 110666.

Dal Santo, N., S. Deparis, and L. Pegolotti, 2020, "Data driven approximation of parametrized PDEs by reduced basis and neural networks," J. Comput. Phys. **416**, 109550.

Demol, P., T. Duguet, A. Ekström, M. Frosini, K. Hebeler, S. König, D. Lee, A. Schwenk, V. Somà, and A. Tichai, 2020, "Improved many-body expansions from eigenvector continuation," Phys. Rev. C **101**, 041302(R).

Demol, P., M. Frosini, A. Tichai, V. Somà, and T. Duguet, 2021, "Bogoliubov many-body perturbation theory under constraint," Ann. Phys. (Amsterdam) **424**, 168358.

de Toledo Piza, A. F. R., E. J. V. de Passos, D. Galetti, M. C. Nemes, and M. M. Watanabe, 1977, "Properties of Griffin-Hill-Wheeler spaces," Phys. Rev. C **15**, 1477.

Djärv, T., A. Ekström, C. Forssén, and H. T. Johansson, 2022, "Bayesian predictions for $A = 6$ nuclei using eigenvector continuation emulators," Phys. Rev. C **105**, 014005.

Drischler, C., J. A. Melendez, R. J. Furnstahl, A. J. Garcia, and X. Zhang, 2023, "BUQEYE guide to projection-based emulators in nuclear physics," Front. Phys. **10**, 92931.

Drischler, C., M. Quinonez, P. Giuliani, A. Lovell, and F. Nunes, 2021, "Toward emulating nuclear reactions using eigenvector continuation," Phys. Lett. B **823**, 136777.

Ekström, A., C. Forssén, G. Hagen, G. R. Jansen, W. Jiang, and T. Papenbrock, 2023a, "What is *ab initio* in nuclear theory?," Front. Phys. **11**, 29094.

Ekström, A., C. Forssén, G. Hagen, G. R. Jansen, T. Papenbrock, and Z. H. Sun, 2023b, "How chiral forces shape neutron-rich Ne and Mg nuclei," arXiv:2305.06955.

Ekström, A., and G. Hagen, 2019, "Global Sensitivity Analysis of Bulk Properties of an Atomic Nucleus," Phys. Rev. Lett. **123**, 252501.

Elhatisari, S., *et al.*, 2016, "Nuclear Binding near a Quantum Phase Transition," Phys. Rev. Lett. **117**, 132501.

Frame, D., R. He, I. Ipsen, D. Lee, D. Lee, and E. Rrapaj, 2018, "Eigenvector Continuation with Subspace Learning," Phys. Rev. Lett. **121**, 032501.

Frame, D. K., 2019, "*Ab initio* simulations of light nuclear systems using eigenvector continuation and auxiliary field Monte Carlo," Ph.D. thesis (Michigan State University).

Francis, A., A. A. Agrawal, J. H. Howard, E. Kökcü, and A. F. Kemper, 2022, "Subspace diagonalization on quantum computers using eigenvector continuation," arXiv:2209.10571.

Franzke, M. C., A. Tichai, K. Hebeler, and A. Schwenk, 2022, "Excited states from eigenvector continuation: The anharmonic oscillator," Phys. Lett. B **830**, 137101.

Franzke, M. C., A. Tichai, K. Hebeler, and A. Schwenk, 2024, "Eigenvector continuation for the pairing Hamiltonian," Phys. Rev. C **109**, 024311.

Fresca, S., A. Manzoni, L. Dedè, and A. Quarteroni, 2020, "Deep learning-based reduced order models in cardiac electrophysiology," PLoS One **15**, e0239416.

Frosini, M., T. Duguet, J.-P. Ebran, B. Bally, H. Hergert, T. R. Rodríguez, R. Roth, J. Yao, and V. Somà, 2022, "Multi-reference many-body perturbation theory for nuclei: III. *Ab initio* calculations at second order in PGCM-PT," Eur. Phys. J. A **58**, 64.

Frosini, M., T. Duguet, J.-P. Ebran, and V. Somà, 2022, "Multi-reference many-body perturbation theory for nuclei: I. Novel PGCM-PT formalism," Eur. Phys. J. A **58**, 62.

Fumagalli, Ivan, Manzoni Andrea, Parolini Nicola, and Verani Marco, 2016, "Reduced basis approximation and *a posteriori* error estimates for parametrized elliptic eigenvalue problems," ESAIM Math. Modell. Numer. Anal. **50**, 1857–1885.

Furnstahl, R. J., A. J. Garcia, P. J. Millican, and X. Zhang, 2020, "Efficient emulators for scattering using eigenvector continuation," Phys. Lett. B **809**, 135719.

Garcia, A. J., C. Drischler, R. J. Furnstahl, J. A. Melendez, and X. Zhang, 2023, "Wave-function-based emulation for nucleon-nucleon scattering in momentum space," Phys. Rev. C **107**, 054001.

Giuliani, P., K. Godbey, E. Bonilla, F. Viens, and J. Piekarewicz, 2023, "Bayes goes fast: Uncertainty quantification for a covariant energy density functional emulated by the reduced basis method," Front. Phys. **10**, 54524.

Griffin, J. J., and J. A. Wheeler, 1957, "Collective motions in nuclei by the method of generator coordinates," Phys. Rev. **108**, 311–327.

Hagen, G., T. Papenbrock, M. Hjorth-Jensen, and D. J. Dean, 2014, "Coupled-cluster computations of atomic nuclei," Rep. Prog. Phys. **77**, 096302.

Herbst, M. F., B. Stamm, S. Wessel, and M. Rizzi, 2022, "Surrogate models for quantum spin systems based on reduced-order modeling," Phys. Rev. E **105**, 045303.

Hergert, H., 2020, "A guided tour of *ab initio* nuclear many-body theory," Front. Phys. **8**, 379.

Hesthaven, J., G. Rozza, and B. Stamm, 2016, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, SpringerBriefs in Mathematics (Springer, New York).

Hicks, C., and D. Lee, 2023, "Trimmed sampling algorithm for the noisy generalized eigenvalue problem," Phys. Rev. Res. **5**, L022001.

Horger, Thomas, Barbara Wohlmuth, and Thomas Dickopf, 2017, "Simultaneous reduced basis approximation of parameterized elliptic eigenvalue problems," ESAIM Math. Modell. Numer. Anal. **51**, 443–465.

Hu, B., *et al.*, 2022, "*Ab initio* predictions link the neutron skin of ^{208}Pb to nuclear forces," Nat. Phys. **18**, 1196–1200.

Ito, K., and S. S. Ravindran, 2001, "Reduced basis method for optimal control of unsteady viscous flows," Int. J. Comput. Fluid Dyn. **15**, 97–113.

Jiang, W. G., C. Forssén, T. Djärv, and G. Hagen, 2024a, "Nuclear-matter saturation and symmetry energy within Δ-full chiral effective field theory," Phys. Rev. C **109**, L061302.

Jiang, W. G., C. Forssén, T. Djärv, and G. Hagen, 2024b, "Emulating *ab initio* computations of infinite nucleonic matter," Phys. Rev. C **109**, 064314.

Kato, T., 2013, *Perturbation Theory for Linear Operators*, Classics in Mathematics Vol. 132 (Springer, New York).

Klos, P., S. König, H.-W. Hammer, J. E. Lynn, and A. Schwenk, 2018, "Signatures of few-body resonances in finite volume," Phys. Rev. C **98**, 034004.

Kohn, W., 1948, "Variational methods in nuclear collision problems," Phys. Rev. **74**, 1763.

Kolmogoroff, A., 1936, "Uber die beste annaherung von funktionen einer gegebenen funktionenklasse [On the best approximation of functions for a given class of functions]," Ann. Math. **37**, 107–110.

Kondo, Y., *et al.*, 2023, "First observation of ^{28}O," Nature (London) **620**, 965–970.

König, S., A. Ekström, K. Hebeler, D. Lee, and A. Schwenk, 2020, "Eigenvector continuation as an efficient and accurate emulator for uncertainty quantification," Phys. Lett. B **810**, 135814.

Lüscher, M., 1991, "Signatures of unstable particles in finite volume," Nucl. Phys. **B364**, 237–251.

Machiels, L., Y. Maday, I. Oliveira, A. Patera, and D. Rovas, 2000, "Output bounds for reduced-basis approximations of symmetric positive definite eigenvalue problems," C. R. Acad. Sci., Ser. I: Math. **331**, 153–158.

McKay, M. D., R. J. Beckman, and W. J. Conover, 1979, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," Technometrics **21**, 239–245.

Mejuto-Zaera, C., and A. F. Kemper, 2023, "Quantum eigenvector continuation for chemistry applications," Electron. Struct. **5**, 045007.

Melendez, J. A., C. Drischler, R. J. Furnstahl, A. J. Garcia, and X. Zhang, 2022, "Model reduction methods for nuclear emulators," J. Phys. G **49**, 102001.

Melendez, J. A., C. Drischler, A. J. Garcia, R. J. Furnstahl, and X. Zhang, 2021, "Fast and accurate emulation of two-body scattering observables without wave functions," Phys. Lett. B **821**, 136608.

Moiseyev, N., 2011, *Non-Hermitian Quantum Mechanics* (Cambridge University Press, Cambridge, England)..

Moiseyev, N., P. Certain, and F. Weinhold, 1978, "Resonance properties of complex-rotated Hamiltonians," Mol. Phys. **36**, 1613–1630.

Nair, P. B., A. J. Keane, and R. S. Langley, 1998, "Improved first-order approximation of eigenvalues and eigenvectors," AIAA J. **36**, 1721–1727.

Newton, R. G., 2002, *Scattering Theory of Waves and Particles* (Dover, Mineola, NY).

Odell, D., P. Giuliani, K. Beyer, M. Catacora-Rios, M. Y. H. Chan, E. Bonilla, R. J. Furnstahl, K. Godbey, and F. M. Nunes, 2024, "ROSE: A reduced-order scattering emulator for optical models," Phys. Rev. C **109**, 044612.

Otsuka, T., M. Honma, T. Mizusaki, N. Shimizu, and Y. Utsuno, 2001, "Monte Carlo shell model for atomic nuclei," Prog. Part. Nucl. Phys. **47**, 319–400.

Pichi, F., A. Quaini, and G. Rozza, 2020, "A reduced order modeling technique to study bifurcating phenomena: Application to the Gross-Pitaevskii equation," SIAM J. Sci. Comput. **42**, B1115–B1135.

Pinkus, A., 1985, *n-Widths in Approximation Theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete Vol. 7 (Springer-Verlag, Berlin).

Quarteroni, A., A. Manzoni, and F. Negri, 2016, *Reduced Basis Methods for Partial Differential Equations: An Introduction*, La Matematica per il 3+2 Vol. 92 (Springer, New York).

Ring, P., and P. Schuck, 1980, *The Nuclear Many-Body Problem* (Springer, Berlin).

Rozza, G., 2005, "Shape design by optimal flow control and reduced basis techniques: Applications to bypass configurations in haemodynamics," Ph.D. thesis (École Polytechnique Fédérale de Lausanne).

Rozza, G., and K. Veroy, 2007, "On the stability of the reduced basis method for Stokes equations in parametrized domains," Comput. Methods Appl. Mech. Eng. **196**, 1244–1260.

Rummukainen, K., and S. Gottlieb, 1995, "Resonance scattering phase shifts on a non-rest-frame lattice," Nucl. Phys. **B450**, 397–436.

Saad, Y., 2011, *Numerical Methods for Large Eigenvalue Problems*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia).

Sarkar, A., and D. Lee, 2021, "Convergence of Eigenvector Continuation," Phys. Rev. Lett. **126**, 032501.

Sarkar, A., and D. Lee, 2022, "Self-learning emulators and eigenvector continuation," Phys. Rev. Res. **4**, 023214.

Sarkar, A., D. Lee, and U.-G. Meißner, 2023, "Floating Block Method for Quantum Monte Carlo Simulations," Phys. Rev. Lett. **131**, 242503.

Shavitt, I., and R. J. Bartlett, 2009, *Many-Body Methods in Chemistry and Physics: MBPT and Coupled-Cluster Theory*, Cambridge Molecular Science (Cambridge University Press, Cambridge, England).

Shimizu, N., T. Abe, Y. Tsunoda, Y. Utsuno, T. Yoshia, T. Mizusaki, M. Honma, and T. Otsuka, 2012, "New-generation Monte Carlo shell model for the K computer era," Prog. Theor. Exp. Phys. 01A205.

Sowiński, T., and M. A. Garcia-March, 2022, "Fundamental limitations of the eigenvalue continuation approach," Phys. Rev. C **106**, 024002.

Svensson, I., A. Ekström, and C. Forssén, 2024, "Inference of the low-energy constants in Δ-full chiral effective field theory including a correlated truncation error," Phys. Rev. C **109**, 064003.

Tews, I., *et al.*, 2022, "Nuclear forces for precision nuclear physics: A collection of perspectives," Few-Body Syst. **63**, 67.

Thom, A. J. W., and M. Head-Gordon, 2009, "Hartree-Fock solutions as a quasidiabatic basis for nonorthogonal configuration interaction," J. Chem. Phys. **131**, 124113.

Tichai, A., J. Langhammer, S. Binder, and R. Roth, 2016, "Hartree-Fock many-body perturbation theory for nuclear ground-states," Phys. Lett. B **756**, 283–288.

Tichai, A., R. Roth, and T. Duguet, 2020, "Many-body perturbation theories for finite nuclei," Front. Phys. **8**, 164.

Tikhomirov, V. M., 1960, "Diameters of sets in function spaces and the theory of best approximations," Russ. Math. Surv. **15**, 75.

Tikhonov, A. N., 1943, "On the stability of inverse problems," Dokl. Akad. Nauk SSSR **39**, 195–198.

Wesolowski, S., I. Svensson, A. Ekström, C. Forssén, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, 2021, "Rigorous constraints on three-nucleon forces in chiral effective field theory from fast and accurate calculations of few-body observables," Phys. Rev. C **104**, 064001.

Wiese, U.-J., 1989, "Identification of resonance parameters from the finite volume energy spectrum," Nucl. Phys. B, Proc. Suppl. **9**, 609–613.

Yapa, N., K. Fossez, and S. König, 2023, "Eigenvector continuation for emulating and extrapolating two-body resonances," Phys. Rev. C **107**, 064316.

Yapa, N., and S. König, 2022, "Volume extrapolation via eigenvector continuation," Phys. Rev. C **106**, 014309.

Yoshida, S., and N. Shimizu, 2022, "Constructing approximate shell-model wavefunctions by eigenvector continuation," Prog. Theor. Exp. Phys. 053D02.

Zhang, X., and R. J. Furnstahl, 2022, "Fast emulation of quantum three-body scattering," Phys. Rev. C **105**, 064004.

Zienkiewicz, O. C., R. L. Taylor, and J. Z. Zhu, 2013, *The Finite Element Method: Its Basis and Fundamentals*, 7th ed. (Butterworth-Heinemann, Oxford).